SEMANTIC AND METAETHICAL PUZZLES ABOUT NORMATIVE LANGUAGE

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Yuna Won

August 2018

SEMANTIC AND METAETHICAL PUZZLES ABOUT NORMATIVE LANGUAGE

Yuna Won, Ph.D.

Cornell University 2018

My three projects here explore some semantic and metaethical problems that are unique to normative language and our normative reasoning. Ch.1 argues that the notion of a contrary-to-duty (CTD) obligation and its role in normative discourse and reasoning are not adequately captured in the standard semantics for *ought*-statements, developed by Angelika Kratzer and David Lewis. I show this by presenting a new puzzle, the CTD Trilemma, using a famous example from Chisholm's Paradox. I claim that two different roles played by *ought*-statements in normative discourse and reasoning have to be captured in any satisfactory semantic account for normative language. I draw the distinction between axiological and deontological *ought*-statements in my version of deontic update semantics and how this distinction nicely solves the CTD Trilemma.

Ch.2 discusses the challenge that the possibility of moral dilemmas poses to the ordering semantic account of *ought*-statements. I critically discuss and compare the advantages and disadvantages of two ordering semantic analyses of *ought*-statements when it comes to moral dilemmas: the conflict and disjunctive accounts. I conclude that two kinds of *ought*-statements defined by the two accounts have to be captured in a single semantic framework in order to adequately describe one's normative predicament and reasoning in a dilemma situation. And, I show one way to do this in the ordering semantic framework.

Ch.3 critically evaluates the real force of the Disagreement-Based Argument

(DBA) in metaethics, exemplified by R.M. Hare's cannibals argument and Horgan and Timmons' Moral Twin Earth argument. I analyze the structure of the DBA and defuse its direct threat to descriptivism by disproving its most problematic assumption: that genuine disagreement between two parties requires that the parties agree in what they mean by the key terms in the dispute. In addition, I maintain that the proponents of descriptivism can and should explain away the elicited intuition used in the DBA: that two causally isolated linguistic communities can have genuine moral disagreements. I show how to do this by using the notion of disagreement about what to do and how to live.

BIOGRAPHICAL SKETCH

Yuna Won grew up in Korea and earned a BA and a MA in Philosophy at Yonsei University. She spent one year at Syracuse University as a visiting student. After that, she entered the Ph.D program at Cornell University, where she worked under Will Starr, Matti Eklund, Janice Dowell and Julia Markovits. She earns her doctoral degree in Philosophy at Cornell in August 2018.

ACKNOWLEDGEMENTS

I would like to express my special appreciation and thanks to my advisor, Prof. Will Starr, and the rest of my special committee members for their guidance and encouragement that made this work possible: Prof. Matti Eklund, Prof. Janice Dowell and Prof. Julia Markovits.

I am grateful that I got to meet many wonderful people and lifelong friends in the department and thank them for standing by me as good friends and fellow philosophers, especially Eric Epstein for his endless encouragement and help during the last summer at Cornell. Special thanks to Prof. Karen Bennett for her support and advice as a placement director. My sincere thanks to the office staff in the department for their thoughtful supports.

My biggest thanks must go to my parents and sister for their love and faith in me over all these years.

TABLE OF CONTENTS

CHAPTER 1

CHISHOLM'S PARADOX REVISITED

1.1 Introduction

A contrary-to-duty (CTD) obligation is a type of conditional obligation that tells us what to do when a primary obligation is violated. Consider the following pair of sentences.

(p1) You ought to keep your promise.

(p2) If you do not keep your promise, you ought to apologize.

(p2) expresses a CTD obligation, when paired with the non-conditional obligation expressed by (p1). This CTD obligation generates a new obligation to apologize for a person who has violated the primary obligation expressed by (p1). CTD obligations are not only used routinely by ordinary speakers, but also indispensable across moral, social and legal discourses. This familiar concept gives rise to one of the most difficult puzzles for Standard Deontic Logic, called *Chisholm's Paradox*. Here's [17]'s original example.

The Chisholm Set

(c1) It ought to be that Jones goes to help his neighbors.

(c2) It ought to be that if he does go he tells them he is coming.

(c3) If he does not go then he ought not to tell them he is coming.

(c4) He does not go.

Although this set of sentences sounds stilted, we have no problem in imagining Jones' normative predicament in which all of these four sentences hold. In Standard Deontic Logic, however, the Chisholm set is either inconsistent or informationally redundant. In particular, the Chisholm set is inconsistent because the following two inferences are valid in Standard Deontic Logic: (c3) and (c4) entail (c5) by Factual Detachment, and (c1) and (c2) entail (c6) by Deontic Detachment.

(c5) Jones ought not to tell his neighbors he is coming.

(c6) Jones ought to tell his neighbors he is coming.

Apparently, (c5) and (c6) express two incompatible obligations. This is why Chisholm's Paradox is often interpreted as a puzzle revealing the difficulty of having both detachment rules for conditional obligations, and motivates many attempts to devise a formal system that validates only one detachment rule, or validates both without the paradoxical results. Thus, it has been widely believed that ordering semantics, which does not validate both detachment rules, is free from Chisholm's Paradox, and thus Chisholm's Paradox has not been discussed much in the ordering semantic framework.

In this paper, I challenge this mainstream diagnosis and widely held belief about ordering semantics. Given the Chisholm set, the most natural interpretation of Jones' normative predicament is this: he violates his primary duty to go help his neighbors, since he does not go, and thus he ought not to tell them he is coming. By proposing a new puzzle, *the CTD Trilemma*, I will show that ordering semantics cannot adequately capture Jones' normative predicament. The CTD Trilemma highlights three characteristic features of the situation where a CTD obligation takes effect: the existence of a violation of duty, the derivation of a new

ought-statement upon the violation of duty, and the absence of pragmatic tension among duties and obligations. It will turn out that ordering semantics is free of the classic puzzle, Chisholm's Paradox, but not of the CTD Trilemma. Thus the notion of CTD obligations still poses a serious challenge to ordering semantics.

I will also show that providing adequate formal representations of two different meanings of *ought*-statements is the key to solving the CTD Trilemma, as well as the first step to developing an adequate formal semantics for normative language. Through this critical discussion I motivate and develop a version of dynamic semantics, which extends ordering semantics based on Veltman([96])'s Update Semantics framework. Finally, I show how my dynamic system solves the CTD Trilemma and explains why the different readings of Chisholm's example are possible.

In sum, this paper aims at making the following contributions:

(1) To appreciate the real challenge raised by CTD obligations by describing a new puzzle, the CTD Trilemma;

(2) To critically evaluate the ordering semantic account of deontic *ought*-statements and conditional obligations by representing the sentences that generate the CTD Trilemma in the ordering semantic framework;

(3) To recognize two types of normativity that *ought*-statements express, and, by that token, the dual nature of our normative language: the descriptive and prescriptive uses of *ought*-statements;

(4) To develop a new dynamic semantics that adequately captures different semantic functions played by *ought*-statements in normative discourse and reasoning, and thereby to solve the CTD Trilemma.

In the next two sections, I briefly introduce Chisholm's Paradox for Standard Deontic Logic, and show why ordering semantics is free from this classic puzzle. In section 4, I discuss the three characteristic features of our reasoning with CTD obligations, and present the CTD Trilemma for ordering semantics. Section 5 draws the distinction between axiological and deontological *ought*-statements in our normative discourse, and discusses why we need two kinds of *ought*-statements to solve the CTD Trilemma. Section 6 develops a version of dynamic semantics that formally incorporats this axiological-deontological distinction. I illustrate how the dynamic account works and how it solves the CTD Trilemma.

1.2 Standard Deontic Logic (SDL) and Chisholm's Paradox

Chisholm's Paradox is one of the most serious challenges to Standard Deontic Logic (SDL). SDL is a normal modal logic with a deontic necessity operator '**OB**' for obligations, and axiomatized as follows.

A1	All tautologies	
A2	$\mathbf{OB}(\phi \supset \psi) \supset (\mathbf{OB}\phi \supset \mathbf{OB}\psi)$	(**OB**-K)
A3	$\mathbf{OB}\phi \supset \neg\mathbf{OB}\neg\phi$	(**OB**-Dual)
R1	If $\vdash \phi$ and $\vdash \phi \supset \psi$ then $\vdash \psi$.	(MP)
R2	If $\vdash \phi$ then $\vdash \mathbf{OB}\phi$.	(**OB**-Nec)

Since in SDL ordinary *if*-statements are expressed by material conditionals, there are two possible combinations of a material conditional and a deontic necessity operator to express conditional obligations: "$\mathbf{OB}(\phi \supset \psi)$" and "$\phi \supset \mathbf{OB}(\psi)$." Therefore, the most straightforward way of representing Chisholm's example is

the following.

The Chisholm Set

(c1)	It ought to be that Jones goes to help his neighbors.	$\mathbf{OB}(h)$
(c2)	It ought to be that if he does go he tells them he is coming.	$\mathbf{OB}(h \supset t)$
(c3)	If he does not go then he ought not to tell them he is coming.	$\neg h \supset \mathbf{OB}(\neg t)$
(c4)	He does not go.	$\neg h$

We can easily imagine the situation where Jones has duties and obligations expressed by (c1)-(c3), but fails to meet his primary obligation ((c1)). In such a case, all four statements of the Chisholm set are true, and each provides independent information about his situation. However, this formal representation of the Chisholm set in SDL leads to a contradiction. Here is a simple proof. (c3) and (c4) entail (c5) by (MP), and (c1) and (c2) entail (c6) by (**OB**-K) and (MP).

(c5)	Jones ought not to tell them he is coming.	$\mathbf{OB}(\neg t)$
(c6)	Jones ought to tell them he is coming.	$\mathbf{OB}(t)$

Since (**OB**-Dual) holds, (c6) is equivalent to $\neg\mathbf{OB}(\neg t)$, which contradicts (c5). Therefore, (c1)-(c4) is *inconsistent*, in the sense that Jones has two mutually incompatible obligations.

Even if (c2) and (c3) are represented as "$h \supset \mathbf{OB}(t)$" and "$\mathbf{OB}(\neg h \supset \neg t)$" respectively, SDL does not faithfully represent the Chisholm set. Since $\phi \supset \psi$ is equivalent to $\neg\phi \vee \psi$, and $\neg\phi$ entails $\phi \supset \psi$, (c4)$\neg h$ entails (c2')$h \supset \mathbf{OB}(t)$, and (c1)$\mathbf{OB}(h)$ entails (c3')$\mathbf{OB}(\neg h \supset \neg t)$. In these cases, the Chisholm set turns out to be *informationally redundant* in that one of the conditional obligations is

entailed by the other non-conditional statement in the Chisholm set. As we can see in the following table, each possible combination of a material conditional and a deontic modal operator in SDL is either inconsistent or informationally redundant.

A	B	C	D
(c1)$\mathbf{OB}(h)$ (c2)$\mathbf{OB}(h \supset t)$ (c3)$\neg h \supset \mathbf{OB}(\neg t)$ (c4)$\neg h$	(c1)$\mathbf{OB}(h)$ (c2') $h \supset \mathbf{OB}(t)$ (c3) $\neg h \supset \mathbf{OB}(\neg t)$ (c4) $\neg h$	(c1)$\mathbf{OB}(h)$ (c2)$\mathbf{OB}(h \supset t)$ (c3')$\mathbf{OB}(\neg h \supset \neg t)$ (c4)$\neg h$	(c1)$\mathbf{OB}(h)$ (c2')$h \supset \mathbf{OB}(t)$ (c3')$\mathbf{OB}(\neg h \supset \neg t)$ (c4)$\neg h$
Inconsistent: (c1) and (c2) entails $\mathbf{OB}(t)$; (c3) and (c4) entails $\mathbf{OB}(\neg t)$.	**Informationally Redundant**: (c4) entails (c2').	**Informationally Redundant**: (c1) entails (c3').	**Informationally Redundant**: (c1) entails (c3'); (c4) entails (c2').

Here's the paradoxical result with Chisholm's example in SDL.

> **Chisholm's Paradox (for Standard Deontic Logic)**
> Intuitively, the Chisholm set is consistent, and each statement in the set tells us independent information. However, the Chisholm set is either inconsistent or informationally redundant in Standard Deontic Logic.

The derivation of a new obligation from a conditional obligation and a factual statement, such as the inference of (c5) from (c3) and (c4), is called *Factual Detachment*. On the other hand, the derivation of a new obligation from a conditional obligation and a non-conditional obligation, such as the derivation of (c6)

from (c1) and (c2), is called *Deontic Detachment*. Many proposed solutions of Chisholm's Paradox focus on resolving the tension between these two detachment rules to avoid inconsistency.[1] However, I will challenge this widely accepted interpretation of Chisholm's Paradox by showing that ordering semantics that does not validate Factual Detachment still fails to adequately represent Jones' normative predicament given the Chisholm set.

1.3 Ordering Semantics and Chisholm's Paradox

In this section, I will briefly introduce the basic ideas of ordering semantics and show why the ordering semantic account for deontic *ought*-statements is free from Chisholm's Paradox. However, we will see that in ordering semantics the Chisholm set entails a counterintuitive prediction about what Jones ought to do when he violates his duty to help his neighbors: even when Jones does not go Jones ought to tell his neighbors he is coming ($\mathbf{O}(t)\&\neg h$). Although ordering semantics can provide more plausible interpretations of Jones' predicament given the Chisholm set, neither of these interpretations captures the concept that is essential for understanding the notion of CTD obligations: the notion of *violation*.

[1] Åqvist([2]) calls this the dilemma on commitment and detachment. This interpretation of Chisholm's Paradox motivates the attempts to devise deontic logic systems that validate Factual Detachment without paradoxical results. Al-Hibri([1]) and Mott([70]) independently develop formal systems in which Factual Detachment is valid while Deontic Detachment is not. Loewer and Belzer([63]) propose a system with both detachment rules without contradiction by introducing temporal elements. Straßer([89]) develops an adaptive logic validating both detachment rules by introducing different notions of obligations. Recently, Willer([100])'s dynamic solution to Chisholm's Paradox is also designed to resolve this tension between the two detachment rules.

1.3.1 The Ordering Semantic Account of *If* and *Ought*

Ordering semantics is a broad family of semantic accounts that interpret conditionals and modals by using quantification over possible worlds and orderings over the possible worlds. Ordering semantics is now an orthodox semantics in linguistics and philosophy of language, most famously developed by Angelika Kratzer and David Lewis.[2] The basic idea of ordering semantics is to understand moral claims as descriptions of the possibilities that are relevant in a given conversational context, and of how these possibilities are ranked relative to a given purpose.

In Kratzer's ordering semantics, the set of accessible possible worlds and a preference ordering on them are determined by two contextual parameters: a *modal base*, a function f mapping an index world i to a set of propositions, $f(i)$, and an *ordering source*, a function mapping i to a set of propositions, $g(i)$, that determines which possible worlds are regarded as ideal as far as that conversational context is concerned. Conditionals in ordering semantics are also modal claims. An *if*-clause plays the role of restricting the set of accessible possible worlds. A deontic modal claim like an *ought*-statement describes what is the case in the deontologically best possible worlds among the accessible possible worlds. Naturally, a conditional *ought*-statement describes what is the case in the deontically ideal worlds in the set of possible worlds that is restricted by the conditional's antecedent.[3]

[2]Different versions of ordering semantics have independently been developed by many linguists and philosophers to deal with different flavors of modals and conditionals: Kratzer([55], [56],[57]), Lewis([60], [61],[62]), Stalnaker([87]), von Wright([98], [40], [41]), Hansson([35]) and Føllesdal & Hilpinen([25]).

[3]Here I assume the standard, simplified ordering semantic account of conditional *ought*-statements. On Kratzer semantics, the semantic function of the antecedent of a conditional is to restrict the domain of the covert modal in the consequent. However, this account is not free of controversy. Alternatively, some argue that a covert epistemic necessity modal takes scope over the deontic modal in order to deal with Kolodny and MacFarlane([54])'s Miner's Paradox ([11]) and the trivial truth of "if ϕ, then Ought(ϕ)" in Kratzer semantics ([15]).

Roughly, "it ought to be that ψ if ϕ ($\mathbf{O}(\psi/\phi)$)" is true if and only if ψ holds in the deontically ideal worlds among accessible ϕ-worlds in the given context.[4] Here is the formal presentation of the truth condition of a conditional obligation ($\mathbf{O}(\psi/\phi)$) in Kratzer semantics.[5]

For a language $\mathcal{L}$ and $\mathcal{M} = \langle \mathcal{W}, f, g, [\![\quad]\!] \rangle$,
$[\![\mathbf{O}(\psi/\phi)]\!]_i^{f,g} = 1$ iff $\{w : w \in \cap f^+(i) \wedge \neg\exists v \in \cap f^+(i) : v \prec_{g(i)} w\} \subseteq [\![\psi]\!]$,
where $\cap f^+(i) = \{w : w \in \cap f(i) \wedge w \in [\![\phi]\!]\}$.[6]

This is read as follows:

A conditional obligation "it ought to be that ψ if ϕ ($\mathbf{O}(\psi/\phi)$)" is true at an index world i if and only if in the subset of ϕ-worlds accessible from i the closest possible worlds to the deontic ideality determined by $g(i)$ are ψ-worlds.

For the sake of convenience, let's use the following simplified gloss:

$\mathbf{O}(\psi/\phi)$ is true in i iff the deontically best worlds among ϕ-worlds accessible from i are ψ-worlds.

[4]Following the convention of dyadic deontic logic, I will represent a conditional *ought*-sentence with a two-place deontic operator '$\mathbf{O}(\psi/\phi)$' and ignore the stylistic difference between "it ought to be that if ϕ then ψ" and "if ϕ then it ought to be that ψ" in a natural language. An unconditional *ought*-statement, '$\mathbf{O}(\phi)$' is an abbreviation for '$\mathbf{O}(\phi/\top)$'—$\top$ stands for a tautology.

[5]For the sake of simplicity, I make the Limit Assumption throughout this paper, according to which there are always some closest possible worlds to the deontic ideal. Without the Limit Assumption, the truth condition of $\mathbf{O}(\psi/\phi)$ is given as follows: $[\![\mathbf{O}(\psi/\phi)]\!]_i^{f,g} = 1$ iff for $u, v, z \in \cap f^+(i), \forall u \exists v [v \preccurlyeq_{g(i)} u \wedge \forall z :$ if $z \preccurlyeq_{g(i)} v$, then $z \in [\![\psi]\!]]$. Here $\cap f^+(i) = \{w : w \in \cap f(i) \wedge w \in [\![\phi]\!]\}$. And $w \preccurlyeq_{g(i)} w'$ means that w is at least as good as w' given $g(i)$; in other words, w meets at least as many propositions in $g(i)$ as w' does.

[6]A model $\mathcal{M}$ for a language $\mathcal{L}$ is a quadruple, $\langle W, f, g, [\![\quad]\!] \rangle$ where $\mathcal{W}$ is a non-empty set of possible worlds, f is a function from the index world, i, to a set of propositions, $\preccurlyeq_{g(i)}$ orders possible worlds given the set of propositions, g(i), and a valuation function $[\![\quad]\!]$ from a well-formed sentence to a set possible worlds; the truth of p at i is represented by $[\![p]\!]_i = 1$. $\cap f(i)$ is a set of possible worlds which satisfy the propositions in $f(i)$. $g(i)$ is a set of the propositions describing the deontic ideality. $w \prec w'$ means that w is strictly better than w'.

1.3.2 Ordering Semantics and Chisholm's Paradox

Given the ordering semantic account of *ought* and *if*, (c1)-(c3) in the Chisholm set have the following truth conditions as deontic modal statements.

(c1) $\mathbf{O}(h)$ is true iff the deontically best worlds are h-worlds.

(c2) $\mathbf{O}(t/h))$ is true iff the deontically best worlds among h-worlds are t-worlds.

(c3) $\mathbf{O}(\neg t/\neg h)$ is true iff the deontically best worlds among $\neg h$-worlds are $\neg t$-worlds.

It is obvious that in ordering semantics each statement of the Chisholm set provides independent information: *no informational redundancy*. More importantly, there is *no inconsistency* in this understanding of the Chisholm set. (c1) and (c2) entail $\mathbf{O}(t)$, whereas (c3) and (c4) do not entail $\mathbf{O}(\neg t)$, because Deontic Detachment is valid, but Factual Detachment is *not valid* in this model.

Deontic Detachment for $\mathbf{O}(\psi/\phi)$: $\mathbf{O}(\psi/\phi) \wedge \mathbf{O}(\phi) \vDash \mathbf{O}(\psi)$

Factual Detachment for $\mathbf{O}(\psi/\phi)$: $\mathbf{O}(\psi/\phi) \wedge \phi \vDash \mathbf{O}(\psi)$

When $\mathbf{O}(\phi)$ and $\mathbf{O}(\psi/\phi)$ are true, $\mathbf{O}(\psi)$ cannot be false in this system. Suppose that the deontically best worlds are ϕ-worlds and among ϕ-worlds the deontically best worlds are ψ-worlds. Then, the deontically best worlds are ψ-worlds as well. Thus, Deontic Detachment, the inference from $\mathbf{O}(\phi)$ and $\mathbf{O}(\psi/\phi)$ to $\mathbf{O}(\psi)$, is valid in ordering semantics. The validity of Deontic Detachment has intuitive appeal. It is natural to infer a new obligation ($\mathbf{O}(\psi)$) from a conditional obligation ($\mathbf{O}(\psi/\phi)$) when its antecedent is obligatory ($\mathbf{O}(\phi)$).

On the other hand, Factual Detachment is not valid for deontic modals. When it comes to deontic modals, it is obvious that an index world, i, is not one of the deontically best possible worlds. For example, in the actual world that we live in, filicide is morally and legally forbidden, and people ought not to commit filicide; however, the actual world is not completely free of it. The actual world is not an ideal world from a moral and legal perspective. Thus it is better not to make the assumption, called *Weak Centering* for deontic modals, according to which the index world i is always one of the best worlds in a relevant sense.[7] Without Weak Centering, the fact that ϕ and $\mathbf{O}(\psi/\phi)$ hold in i does not guarantee that $\mathbf{O}(\psi)$ holds in i—it is possible that $\mathbf{O}(\phi)$ is false in i. Therefore, Factual Detachment is invalid in ordering semantics. As a result, (c5) does not follow from (c3) and (c4), whereas (c6) does follow from (c1) and (c2). There is no inconsistency in the Chisholm set.

(c5) Jones ought not to tell them he is coming. $\mathbf{O}(\neg t)$

(c6) Jones ought to tell them he is coming $\mathbf{O}(t)$

In sum, ordering semantics is free from the paradoxical results of the Chisholm set thatwe saw in SDL, because the antecedent of a conditional is a domain restrictor, and because Factual Detachment is not valid in ordering semantics. Each sentence of the Chisholm set provides independent information, and the set does not entail a contradiction.

[7]Formally, Weak Centering amounts to the following assumption about an ordering source $g(i)$: $i \in \{w : w \in \cap f^+(i) \wedge \neg\exists v \in \cap f^+(i) : v \prec_{g(i)} w\}$.

1.4 The CTD Trilemma : The New Puzzle

We have seen that ordering semantics represents the Chisholm set without inconsistency or redundancy; thus, it is free of Chisholm's Paradox. However, whether this is a faithful representation of our reasoning with the Chisholm set is another question. To answer to this question, I am going to examine three different ways of interpreting the Chisholm set and consider whether each of them can be adequately represented in the ordering semantic framework.

1.4.1 Three Interpretations of the Chisholm Set

The sentences of the Chisholm set call to mind three different types of situation that Jones might be in: the According-to-Duty (ATD) Situation, the Overriding Situation, and the CTD Situation. The first type of situation is where Jones has not gone to help his neighbors yet, but it is still the most desirable thing for him to do and Jones can go help his neighbors. Suppose, for example, that Jones is out of town, but he knows that his weekend will be well-spent if he helps his neighbors move tomorrow. Imagine that he can still help them tomorrow if he catches the train back to town tonight. In this situation, what he ought to do is to go help his neighbors and let them know he is coming. So the *ought*-statements, (c1)-(c3), hold, and the derived *ought*-statement, (c6), holds, but not (c5). Let's call this situation the According-to-Duty (ATD) Situation.

Second, we can imagine the situation where it would be a good thing for him to help his neighbors, but it is outside his ability. For example, suppose that Jones is out of town and knows that his neighbors need his help tomorrow, but all trains back to town were cancelled till tomorrow afternoon due to the bad weather. Given

Jones' predicament, he ought not to tell them he is coming, and we cannot say that Jones ought to go help his neighbors, since he cannot. Thus, (c2)-(c5) hold in this situation. Let's call this situation the Overriding Situation.

Third, we can imagine the situation where Jones ought to go help his neighbors on a certain day because he has a duty to help them as a community member, but Jones has failed or will fail to fulfill the duty. In this case, his duty as a community will remain even after his failure. Suppose, for example, that Jones is out of town, and has missed the last train back to town in time. He won't be able to help his neighbors tomorrow, so he ought not to tell them he is coming. But it is still true that he ought to help his neighbors, since it's his duty as a community member. In this situation (c1)-(c5) hold. Let's call this situation the CTD Situation.

Notice that the CTD Situation is different from the other two situations in that (c1) is understood as expressing Jones' duty or obligation in the CTD Situation, whereas in the ATD and Overriding Situations, we read (c1) as a sort of recommendation or advice for Jones, given his current circumstance in the ATD and Overriding Situations. Furthermore, in the ATD Situation Jones can go help his neighbors, whereas in the Overriding Situation he cannot.

The standard ordering semantic account of *ought*s and *if*s can adequately represent the ATD Situation and the Overriding Situation and explain nicely why the truth of (c1) depends on whether Jones can go help his neighbors. Here's how that goes. It is a logical truth that either Jones can go help his neighbors, or Jones cannot go help his neighbors ($\Diamond h \vee \neg\Diamond h$). In ordering semantics, when $\Diamond h$ is true at w, there is at least one h-possible world accessible from w. On the other hand, when $\neg\Diamond h$ is true at w, there is no h-world among all the possible worlds accessible from w. Therefore, the two possible situations are captured by

different modal bases: $\neg h \notin f_1$ and $\neg h \in f_2$, respectively. Suppose that the relevant ordering source is $g = \{h, h \supset t, \neg h \supset \neg t\}$. In the situation where Jones can go ($\Diamond h$), the relevant contextual parameters are f_1 and g. Given f_1 and g, there are some h-worlds in the relevant domain, and h-worlds are the deontically best worlds. Therefore, Jones ought to go help his neighbors and tell them he is coming ($\mathbf{O}(h) \wedge \mathbf{O}(t)$). On the other hand, in the situation where Jones cannot go to help his neighbors ($\neg\Diamond h$), the relevant contextual parameters are f_2 and g. Given f_2 and g, there are no h-worlds in the relevant domain. And, among $\neg h$-worlds, $\neg t$-worlds are better than t-worlds. Therefore, Jones ought not to tell them he is coming ($\mathbf{O}(\neg t)$); but, importantly, it is not the case that he still has an obligation to go help his neighbors ($\neg\mathbf{O}(h)$).[8]

1.4.2 The CTD Situation and The CTD Trilemma

As we just saw, the standard ordering semantics framework successfully represents the ATD Situation and the Overriding Situation, and explains the relation between what Jones ought to do and what Jones can do. But what about the CTD Situation? This is an important question, given that the CTD Situation seems to be the most natural way of understanding Jones' predicament given the Chisholm set. (That is because unlike the ATD situation, Jones is not going to ("does not", read timelessly) go help his neighbors ((c4)); but unlike the Overriding situation, Jones nonetheless has an obligation to do so.) In this section, I will first identify

[8]This type of response to Chisholm's Paradox in the ordering semantic framework is given and briefly discussed in Portner ([73]). Also, this modal solution is closely related to classic temporal solutions to Chisholm's Paradox. The ATD Situation and the Overriding Situation naturally correspond to the two time periods before and after a certain point: Jones can go help his neighbors before t_n and Jones cannot go to help his neighbors after t_n. Thus, several temporal solutions to Chisholm's Paradox basically understand Jones' predicament as either the ATD Situation and the Overriding Situation in the modal solution here. For such temporal approaches to solving Chisholm's Paradox, see Feldman ([23]) and Loewer and Belzer ([63]).

three characteristic features of the CTD Situation that explain why the Extended Chisholm set, (c1)-(c5), holds for Jones in the CTD Situation. Then, I will show how the CTD Situation generates a new CTD puzzle.

The first feature of the CTD Situation is that in it, (c1) is understood as Jones' *duty*, not merely a good thing to do. In the CTD Situation, Jones' duty as a community member, expressed by (c1), does not disappear simply because he cannot fulfill it in a certain circumstance. Rather, the duty stands unless there is a good reason to exempt him. Therefore, in this situation, when Jones does not go help his neighbors, he violates his primary duty to them. The violation of a duty or obligation can be formally expressed by the conjunction "$\mathbf{O}(\phi) \wedge \neg\phi$" in ordering semantics. So (c1) and (c4) hold in the CTD Situation.

Second, Jones' CTD obligation ((c3)) tells him what to do upon the violation of his duty expressed by (c1). A newly-derived *ought*-statement, (c5), guides Jones concerning what to do in the CTD Situation. Telling one what to do upon the violation of a primary duty is the main function of a CTD obligation. Therefore, the inference from a violation of a primary duty to there being a CTD obligation has to be validated in a formal system, for it to adequately capture the reasoning connected with CTD obligations. Thus (c5) ought to come out true.

Third, the CTD Situation is not an aberrant situation. Admittedly, the CTD situation is not an ideal situation, due to the violation; and indeed there is a sort of tension between Jones' primary duty and what he ought to do in the CTD Situation.[9] However, it is not an aberrant situation, the way that a moral dilemma is. An agent facing a moral dilemma is torn between two equally pressing duties

[9]Prakken and Sergot ([74], [75]) call this tension between a primary obligation and a derlived *ought*-statement in the CTD Situation *Pragmatic Oddity*. Although it is not logically inconsistent, it is absurd to require Jones to go help his neighbors and not to tell them he is coming given the Chisholm set.

or obligations, and the existence of a dilemma is often considered as an indicator of the inconsistency in a set of duties and obligations. By contrast, Jones in the CTD Situation does not experience a dilemma—it is clear what he ought to do. Therefore, there is no good reason to suspect that the Extended Chisholm set is inconsistent.

These three characteristics of the CTD Situation naturally establish the following three theoretical desiderata for the faithful formal representation of the CTD Situation.

I. Violation of a Primary Obligation

In the CTD Situation, the relevant agent's primary obligation holds, but it is violated: $\mathbf{O}(\phi) \wedge \neg\phi$.

II. Inference of a New *Ought*-Statement

In the CTD Situation, there is a newly derived *ought*-statement that tells the agent what to do given that he or she violates the primary obligation: $\mathbf{O}(\psi)$ from $\neg\phi$ and $\mathbf{O}(\psi/\neg\phi)$.

III. Compatibility

In the CTD Situation, there is no inconsistency among the *ought*-statements expressing the agent's duties and obligations.

Meeting all three desiderata amounts to representing the following extended Chisholm set without inconsistency in ordering semantics.

The Extended Chisholm Set

(c1)	It ought to be that Jones goes to help his neighbors.	$\mathbf{O}(h)$
(c2)	It ought to be that if he does go he tells them he is coming.	$\mathbf{O}(t/h)$
(c3)	If he does not go then he ought not to tell them he is coming.	$\mathbf{O}(\neg t/\neg h)$
(c4)	He does not go.	$\neg h$
(c5)	He ought not to tell them he is coming.	$\mathbf{O}(\neg t)$

However, the standard ordering semantic account of *ought*s and *if* faces difficulties in meeting the three desiderata and representing the Extended Chisholm set without inconsistency. There are a couple of possible ordering semantic approaches, but each of them only meets two desiderata rather than all three at the same time.

First, a version of ordering semantics that rejects Weak Centering can easily express the violation of Jones' primary obligation by using the conjunction $\mathbf{O}(h) \wedge \neg h$. The rejection of Weak Centering means that the actual world (or an index world) is not always one of the ideal possible worlds; thus, $\mathbf{O}(h)$ could be true in the actual world, where h does not hold. However, the rejection of Weak Centering results in the invalidity of Factual Detachment. Thus, this version of ordering semantics cannot explain why it is natural to infer the new *ought*-statement, (c5) $\mathbf{O}(\neg t)$, when the primary obligation is violated ($\mathbf{O}(h) \wedge \neg h$), and his CTD obligation ($\mathbf{O}(\neg t/\neg h)$) holds. So, on this approach, ordering semantics meets the violation requirement (I) but not the new *ought*-statement requirement (II).

Some might wonder whether a version of ordering semantics that assumes Weak Centering for deontic modals can avoid the CTD Trilemma. It is true that the invalidity of Factual Detachment is a huge price to pay, and by assuming Weak

Centering we can keep Factual Detachment for conditional *ought*-statements.[10] Accordingly, let's now consider a version of ordering semantics that assumes Weak Centering and validates Factual Detachment. Once it is assumed that the actual world is one of the ideal possible worlds when evaluating deontic modals, as Weak Centering says, the new *ought*-statement requirement (II) is easily met. If Jones does not go ($\neg h$) and his CTD obligation ($\mathbf{O}(\neg t/\neg h)$) still holds, then indeed he ought not to tell them he is coming ($\mathbf{O}(\neg t)$), because the actual world is one of the deontically ideal possible worlds in which $\neg h$ holds. However, the violation of Jones' primary obligation, expressed by the conjunction $\mathbf{O}(h) \wedge \neg h$, is not possible in this system, since, according to Weak Centering, the truth of $\mathbf{O}(h)$ entails h at i. So, the version of ordering semantics under consideration fails to meet the violation requirement (I). In sum, whether or not Weak Centering is assumed, ordering semantics faces the CTD Trilemma: if Weak Centering is assumed, then it fails to meet the violation requirement (I), and if not, then it fails to meet the new *ought*-statement requirement (II) due to the invalidity of Factual Detachment.

A third approach for ordering semantics is to explain why inferring (c5) from (c3) and (c4) is reasonable without appealing to the validity of Factual Detachment, which assumes Weak Centering. Here is one example of vindicating the inference without assuming Weak Centering. Following [54], one can argue that the inference from $\mathbf{O}(\psi/\phi) \wedge \phi$ to $\mathbf{O}(\psi)$ is not valid, but quasi-valid. So here is [54]'s notion of *quasi-validity*:

> An inference from premise $\phi_1, \phi_2, ..., \phi_n$ to conclusion ψ is *quasi-valid* iff the

[10] It is true that the invalidity of Factual Detachment itself is the serious drawback of rejecting Weak Centering. The invalidity of Factual Detachment is at odds with our intuitive understanding of conditional obligations and everyday uses of them. Factual Detachment is an incredibly useful inference pattern in our normative discourse, allowing us infer what to do in specific situations. As Van Eck ([94]) asks, "How can we take seriously a conditional obligation if it cannot, by way of detachment, lead to an unconditional obligation?" (p.263).

inference from $\Box_e\phi_1, \Box_e\phi_2, ..., \Box_e\phi_n$ to ψ is valid.

This notion of quasi-validity relies on the informal epistemic property of inference, according to which the conclusion must be true when the premises are *known* to be true. And, here is a direct application of quasi-validity to the reasoning with deontic Factual Detachment:

> An inference from premises $\mathbf{O}(\psi/\phi)$ and ϕ to conclusion $\mathbf{O}(\psi)$ is *quasi-valid* iff the inference from $\Box_e[\mathbf{O}(\psi/\phi)]$ and $\Box_e\phi$ to $\mathbf{O}(\psi)$ is valid.

Whether Factual Detachment for $\mathbf{O}(\psi/\phi)$ is quasi-valid or not completely depends on whether the modal base for deontic *ought* is the epistemic modal base or not. If not, then we cannot extend the notion of quasi-validity to deontic reasoning. However, even if so, there are further problems. What is derived from a quasi-valid inference and what is derived from a valid inference are in tension: (c5) $\mathbf{O}(\neg t)$ and (c6) $\mathbf{O}(t)$. This is semantically inconsistent given the dual of $\mathbf{O}(\phi)$ with $\neg\mathbf{O}(\neg\phi)$. But even without the dual, this is not a faithful representation of Jones' predicament, because by hypothesis, Jones does not face a dilemma between two conflicting *ought*-statements in the CTD Situation. So this approach fails to meet the compatibility requirement (III) even though it meets (I) and (II).[11]

As a result, the three characteristic features of the CTD Situation generate the following trilemma for ordering semantics:

The CTD Trilemma

[11] Instead of introducing the notion of quasi-validity, the proponents of Kratzer semantics might take a flexible contextualist approach by using different pairs of contextual parameters $\langle f, g\rangle$ for (c1) and for (c3) to infer (c5). This approach would meet the three desiderata but have the consequence that the Extended Chisholm set does not hold in one single context.

The ordering semantic account of *ought*s and *if* can meet two of the three features of the CTD Situation, but not all three of them. Therefore, it fails to provide a faithful representation of the CTD Situation.

The CTD Trilemma is a new puzzle raised by the CTD Situation, one of the most natural understandings of Jones' predicament given the Chisholm set. The CTD Trilemma shows that merely representing the Chisholm set without inconsistency and redundancy is not enough for a faithful formal representation of natural reasoning with a CTD obligation. Rather, a faithful semantics must also capture the Extended Chisholm set without inconsistency. Moreover, the three requirements of the CTD Trilemma that characterize three features of the CTD Situation are not easily met in ordering semantics. Even when we meet (I) and (II), the direct tension between the two derived *ought*-statements, (c5) and (c6), and the pragmatic oddity between (c1) and (c5) are unavoidable unless we can say that (c5) is a different kind of *ought*-statement from (c1) and (c6). This is precisely what I wish to propose.

1.5 The CTD Trilemma and Two Kinds of *Ought*-Statements

The CTD Trilemma witnesses three desiderata for adequately representing the CTD Situation, and significant difficulties that the most basic version of ordering semantics has in meeting them. Ordering semantics cannot meet the three desiderata unless it embraces the idea that the *ought* in (c1) and the *ought* in (c5) are not the same kind of *ought*. This suggests that the key step toward diagnosing and solving the CTD Trilemma is to examine whether there really are two different kinds of *ought*-statements in play in our reasoning and discourse concerning CTD obli-

gations, and to provide a plausible and general distinction between them. In this section, I will argue that there are indeed two different types of *ought*-statements, and that recognizing this is required for solving the CTD Trilemma. Having two different types of *ought*-statements on hand will enable us to read the Chisholm set in different ways.

1.5.1 Axiological and Deontological *Ought*-Statements

Here I claim that there are two different types of *ought*-statements, relating to two different kinds of normative considerations that operate in our normative discourse and reasoning. Consider an example of a simple *ought*-statement, "Sue ought to keep her promise". This *ought*-statement holds in two different situations: when it is a good thing for Sue to keep her promise in a given circumstance, and when there are certain duties or obligations that require Sue to keep her promise regardless of her current circumstances. When used in the former situation, this *ought*-statement expresses the speaker's evaluative judgment about what is the most desirable thing for Sue to do (all-things-considered) in her current circumstance. When used in the latter situation, the same *ought*-statement expresses what it is required for Sue to do, as dictated by her relevant duties or obligations. Let's call the former kind of *ought*-statements *axiological ought*-statements, and call the latter kind *deontological ought*-statements.

The axiological-deontological distinction relates to two different types of normative reasons, what is good all-things-considered in a given circumstance and what is required by a set of relevant duties and principles, respectively. This distinction in normative reasons is reflected in two distinct patterns of reasoning that we can observe in our reasoning and discourse involving *ought*-statements.

First, the distinction between what is good all-things-considered and what is required by normative principles is indispensable for understanding the notions of supererogation and violation. A certain action is supererogatory when it is good, but not required to perform the action. And, there is a violation when we fail to do what we are required to do. Recall Singer's famous observation in his "Famine, Affluence and Morality ([84])," for example. To give most of my income to charity and save children dying from hunger and curable diseases is morally the most desirable thing for me to do, and it is within my capability to do this without downgrading my own life beyond acceptable limits. It is indeed a good thing to do to give most of my income to charity, but it is not required. By not doing so, I do not violate anything. Whether or not Singer is right about our moral duties, the controversy around whether we are morally required to do what is the most morally desirable thing to do reveals an incontestable conceptual distinction between what is good and what is required, and thus the possibility of supererogation.

The second characteristic difference between the two normative reasons is that they are different in their sensitivity to what an agent can and can't do. For example, the axiological $ought_A$-statement, "I ought to give most of my income to charity," stops holding as soon as my income reaches a certain point in the lower end of the income scale. This is because the evaluative judgment about what is good or the most desirable thing to do changes as the circumstance changes. Therefore, axiological *ought*-statements are sensitive to contingent facts and an agent's practical capability. On the other hand, the moral requirements expressed by deontic *ought*-statements are less sensitive to agents' circumstances and capabilities. For example, my civil duty to pay taxes does not disappear unless I stop being recognized as an economic agent. Regardless of how much I make or whether I might be bankrupt after paying my taxes, the deontological *ought*-statement, "I

$ought_D$ to pay taxes," holds. I could be exempt from my duty when I meet certain conditions or lose my status as an economic agent. But my incapability of fulfilling this duty by, e.g., simply not having the money or by missing the due date, does not dissolve my civil duty. Even when I violate this duty, it holds as long as I am a proper member of this society.

Consequently, *ought*-statements that relate to different types of normative reasons behave differently in our normative discourse and reasoning. I'll now show that the axiological-deontological *ought*-distinction based on different normative reasons provides us a non-*ad hoc* explanation of why several different interpretations of the Chisholm set are possible.

1.5.2 The Two-*Ought* Solution

Suppose that (c1)-(c3) in the Chisholm set are true axiological *ought*-statements. This means, Jones has some normative reasons to do the conditional or unconditional actions recommended by (c1)-(c3) because they are good things to do all-things-considered in his circumstance. Jones $ought_A$ to go help his neighbors because it is a good thing for him to do. As long as he can go help his neighbors, (c1) as an axiological *ought*-statement holds. But as soon as it turns out that he cannot go help his neighbors, (c1) as an axiological *ought*-statement stops holding, because in his circumstance it is not a good thing to do anymore. In either case, there is no violation because (c1) says what is good for Jones to do, not what he is required to do. Therefore, when (c1) is read as an axiological *ought*-statement, we naturally understand Jones' predicament as the ATD Situation or the Overriding Situation, depending on whether Jones can or cannot go help his neighbors. In the ATD Situation, the following conjunction of axiological *ought*-statements holds:

> Jones $ought_A$ to go help his neighbors and $ought_A$ to tell them he is coming.

In the Overriding Situation, the following holds:

> Jones $ought_A$ not to tell them he is coming, since he does not go; and, it is not the case that he $ought_A$ to go help is neighbors.

On the other hand, the CTD Situation is the most plausible understanding of Jones' predicament when reading (c1) as a deontological *ought*-statement. Suppose that (c1) is a true deontological *ought*-statement, because it is required for Jones to go help his neighbors by his relevant duty, such as his duty as a community member. And, suppose that Jones has not and is not going to go help his neighbors due to his negligence or laziness, or the occurrence of some unexpected event. His failure to help his neighbors constitutes a violation of his duty as community member, and given this violation the most desirable thing for him to do is not tell his neighbors that he is coming. About this CTD Situation, therefore, we say as follows:

> Jones $ought_D$ to go help his neighbors, but he is not going to. So he $ought_A$ not to tell them he is coming.

Now it is clear how the axiological and deontological *ought* distinction helps us avoid the CTD Trilemma. First, understanding (c1) as a deontological *ought*-statement is necessary to make sense of the violation in the CTD Situation. Second, given that Jones does not go ((c4)), what is most desirable is for him not to tell his neighbors that he is coming. Last, since this newly derived *ought*-statement in Jones' current predicament is an axiological *ought*-statement, it is not in conflict with Jones' duty to go help his neighbors, which is expressed by a deontological

ought-statement. With this distinction in hand, we can express the Extended Chisholm set as follows without inconsistency.

The Extended Chisholm Set in the CTD Situation

(c1) Jones ought$_D$ to go help his neighbors.

(c2) If Jones does go, he ought$_D$ to tell them he is coming.

(c3) If he does not go, he ought$_D$ not to tell them he is coming.

(c4) He does not go.

(c5) He ought$_A$ not to tell them he is coming.

The advantage of introducing two *ought*s here is obvious. With these two distinct kinds of *ought*-statements we can avoid the CTD Trilemma.[12]

So far, I have presented the CTD Trilemma as a problem for ordering semantics. But now that the deontological-axiological distinction emerges as a solution, one might wonder whether advocates of ordering semantics can find a way to accommodate that distinction. To that effect, some proponents of Kratzer semantics might argue that different contextual parameters are assigned to deontological and axiological *ought*-statements in the Extended Chisholm set, and they are determined via speakers' intentions or contextually salient features. Here is a quick sketch of

[12]This two *ought* solution illustrated here is different from the various recent approaches to solve Chisholm's Paradox by validating both Factual and Deontic Detachment in different systems. Willer ([100])'s non-monotonic analysis of deontic *ought* in a dynamic framework. Arregui ([3], [4]) claims that the contrast between *should* and *should have* should be drawn in terms of their differences in context set presuppositions and based on this distinction explains why the two inference patterns in the Chisholm set. Silk ([83]) validates both detachment rules without tensions by understanding *ought*s in a conditional statement as relative to a contextually salient ordering source. My two *ought* solution and the axiological and deontological *ought* distinction are not directly motivated by validating two detachment rules, rather by the conceptual contrast between two *ought*s and between two different normative considerations they reflect in our normative reasoning.

what this contextualist solution would look like. First, this contextualist approach rejects Weak Centering in order to explain the violation of Jones' primary duty when he does not go help his neighbors. Second, it assigns contextual values for f and g that would allow the inference of (c5) from (c3) and (c4). Third, different contextual values for f and g are assigned for evaluating (c1) and (c5) to avoid the tension between them. This means that the value of g remains unchanged, but the different values of f are in play for (c1) and (c5). For example, suppose that $f(w)$ includes some h-worlds where Jones goes help his neighbors, but $f'(w)$ has no h-worlds. While h-worlds are ranked most highly given $\langle f(w), g(w)\rangle$, $\neg t$-worlds are ranked most highly given $\langle f'(w), g(w)\rangle$. So this flexible contextualist approach can represent the Extended Chisholm set as true as follows:

$$[\![\mathbf{O}(h))]\!]^w_{f,g} = [\![\mathbf{O}(t/h))]\!]^w_{f,g} = [\![\mathbf{O}(\neg t/\neg h))]\!]^w_{f,g} = [\![\neg h]\!]^w = [\![\mathbf{O}(h))]\!]^w_{f',g} = 1$$

This contextualist approach basically claims that we consider different domains of possibilities for (c1)-(c3) and for (c5), and that this is the semantic distinction between deontological and axiological *ought*-statements.

However, advocates of this contextualist two-*ought* solution must explain why we need to assign different contextual parameters for *ought*-statements in the Extended Chisholm set. Given that we can articulate the Extended Chisholm set in the same breath, and given that the Extended Chisholm set seems to hold of a single situation (viz., Jones' predicament) with no shift in context, why should we think that there is more than one value in play for any relevant contextual parameter? More importantly, one can question whether it really is adequate to formally represent the distinction between axiological and deontological *ought*-statements in terms of different contextual values assigned to *ought*-statements.

To make sense of this distinction in the ordering semantic framework, one says that an axiological *ought*-statement tells us which of the possibilities restricted by the relevant facts of a certain circumstance is most highly ranked; while a deontological *ought*-statement tells us which possibility in the *unrestricted* domain is most highly ranked. The proponents of Kratzer semantics, in particular contextualists, might be happy with this formal distinction that simply posits two different modal bases for the two types of statement. However, it seems that the conceptual distinction that ordinary normative reasoning makes between what is good and what is required does not simply reduce to necessity claims about different domains of possibilities. That is, in the same breath, we can articulate both a type of obligation that Jones has which is modulated by changes in his situation, and another type of obligation, emanating directly from his status as a neighbor, which continues to hold regardless of such changes.

The ordering semantic account of modals is basically axiological in that it gives the truth condition of an *ought*-statement in terms of comparative betterness between relevant possible worlds. For example, "Sue ought to keep her promise" is true iff the possible worlds in which Sue keeps her promise are better than those in which she does not. So the ordering semantic analysis of *ought*-statements is conceptually based on the evaluative judgment about what worlds are better than others. It does *not* capture the meaning of a deontological *ought*-statement that holds whenever there is a relevant set of duties or obligations, regardless of how that affects the comparative quality of worlds. This worry that I articulate about ordering semantics resonates with van Fraassen's ([95]) criticism of the semantic foundation of deontic logic for its conceptual reduction of (directive) normativity into comparative betterness. He argues that defining the meaning of *ought*-statements in terms of comparative betterness among alternative possibil-

ities which are rated on a certain scale of value commits us to the controversial philosophical claim that the axiological is prior to or explains the deontic—which he calls the *axiological thesis*.[13]

The conceptual distinction between axiological and deontological *ought*-statements helps us see that van Fraassen's concern remains for ordering semantics: ordering semantics does not assign an adequate truth-condition for deontological *ought*-statements: deontological *ought*-statements are sensitive to the existence of normative principles. As we have seen, true deontological *ought*-statements directly reflect what is required by the relevant set of normative principles. Normative principles, such as duties and obligations, shape one's preference ordering over possible actions. Thus, normative principles are often considered as contributing to a relevant ordering source in ordering semantics. However, in ordering semantics these normative principles only indirectly contribute to the truth-condition of *ought*-statements, via an ordering source. So, what is required by one's duty or obligation cannot be expressed by a true *ought*-statement unless it is the most desirable possibility in the contextually relevant domain of possibilities. But the point of the CTD Situation and the CTD Trilemma is that we reason differently when thinking in deontological terms than we do when thinking in axiological terms: a deontological *ought*-statement holds when there exists a corresponding normative principle, even when the relevant agent cannot fulfill it.

To capture the deontological meaning that *ought*-statements can have, I will

[13]He points out that there are some immediate problems with the axiological thesis. First, if the axiological thesis holds, then a moral dilemma situation in which two incompatible obligations hold at the same time is impossible. Second, the axiological reduction as a result of the axiological thesis does not leave room for supererogatory acts, acts that are morally good but not required. This problem is often called "the Problem of Supererogation." These results are highly undesirable, since a semantic theory, conceived as nothing but a formal account of the meanings of sentences in a certain domain, should not make substantial commitments to first order philosophical claims, especially controversial ones. This methodological principle is, at any rate, standardly accepted by philosophical logicians.

propose a version of dynamic semantics in the next section. This dynamic semantics is designed to enable us to keep track of the two different kinds of normative reasons for *ought*-statements, and systematically explain the interaction between two kinds of *ought*s. So this model will fill an important theoretical gap neglected in most versions of ordering semantics, explaining the mechanism of how some *ought*-statements contribute to a contextual parameter such as an ordering source, and how an ordering source changes over time as a discourse evolves and a context changes.

We have seen that the deontological meaning of an *ought*-statement is not adequately captured by the truth-condition of deontic *ought*-statements in ordering semantics. Unlike axiological *ought*-statements, deontological *ought*-statements are about a set of relevant normative principles. In ordering semantics, the relevant normative principles in a context are taken into an ordering source, and they contribute to the preference ordering over relevant possibilities—these being determined by a modal base. Deontological *ought*-statements are about the principles that contribute to an ordering source. So it is natural to understand that deontological *ought*-statements contribute to an ordering source, and they hold when there is a corresponding proposition in the ordering source. The problem is that this connection between true deontological *ought*-statements and an ordering source is not semantically explicit in ordering semantics. Rather, it is left in the realm of pragmatics, since the ordering source is given as a contextual parameter. To remedy this limitation and provide a semantic account of the meaning of deontological *ought*-statements, I propose a dynamic implementation of ordering semantics based on Veltman's Update Semantics.

1.6 Dynamic Update Semantics for Two *Ought*-Statements

Here I put forward a dynamic implementation of the ordering semantics, which is formally inspired by Veltman([96])'s Update Semantics. As a dynamic account, the core idea of Update Semantics is to understand the meaning of a sentence as consisting in how it changes or updates the information states of language users.Therefore, the meaning of a sentence is not only interpreted against a relevant information state or a context, but also in terms of its systematic contribution to relevant informational states or contexts. This is why the dynamic meaning of a sentence is often called a context change potential, a function from one information state or context to another.[14] Veltman nicely contrasts the difference between static and dynamic systems in how they understand meanings.

> The slogan "You know the meaning of a sentence if you know the conditions under which it is true" is replaced by this one: "You know the meaning of a sentence if you know the change it brings about in the information state of anyone who accepts the news conveyed by it" ([96], 1).

An information state reflects what a relevant agent or group knows, and is represented by a set of possible worlds that have not been ruled out. Suppose that A and B are the only participants in a certain conversation, and A says, "it's snowing." The meaning of this utterance is, of course, interpreted against contextual elements like when and where A speaks. Not only that, but also A's utterance

[14]The origin of this dynamic understanding of meaning can be traced back to Stalnaker's work on presupposition and assertion ([86]). This core idea of dynamic semantics is endorsed and developed by many authors to provide a better explanation of various linguistic phenomena: Kamp ([52], [53])'s Discourse Representation Theory, Groenendijk and Stokhof ([34])'s Dynamic Predicate Logic, Heim([39])'s File Change Semantics, Veltman([96])'s Update Semantics, and recent dynamic approaches like Beaver ([6]), von Fintel and Gillies ([97]), Starr ([88]), and Willer([99],[100]).

brings about a systematic change to B's information state or their common information state. By accepting A's utterance, B's information state evolves into a new state that accommodates this new piece of information by eliminating the possibilities incompatible with the new information. For example, after learning that it's snowing, B cannot believe that it is not snowing outside. Therefore, the meaning of the sentence "it's snowing," is the function from one information state to the information state that rules out all the possibilities incompatible with "it's snowing." Let's call this process of incorporating new information and evolving into a new information state *update*. When a certain information state is already compatible with a belief ϕ, let's say ϕ is accepted in the information state—this means, updating this information state with ϕ does not bring about any change in the information state, but returns the very same information state.[15]

In a dynamic account, different sentences have different meanings because they change information states differently or in different ways. So eventually I will argue that axiological and deontological *ought*-statements change information states differently and are accepted by different information states. To capture the difference in their dynamic meanings and the interaction between them, my update semantics for normative language will introduce *structured information state* instead of simple information state represented by a set of possible worlds.[16]

[15]Although here I use this simplified toy example, the notion of information state is a highly abstract and idealized concept in most versions of dynamic semantics. An information state models what is believed or accepted not only by a relevant agent such as an individual, a group or a loosely defined discourse participants, but also by a body of general knowledge or a common ground of a group or society. Also, it is assumed that all utterances are sincere and felicitously made in this paper. Also, this paper sets aside the important issues regarding acceptability and assertability that could be relative to information states.

[16]Torre and Tan ([93]) also put forward Deontic Update Semantics based on Veltman's system. The major difference of theirs from my version is that they do not distinguish between axiological *ought*-statements and deontological *ought*-statements. The dynamic meaning of "it is obligatory that ϕ if ψ" in their system is very similar to the meaning of deontological *ought*s in my version, but it does not account for axiological meaning of *ought*s. As a result, it lacks the story about the interaction between two types of *ought*s in normative reasoning. Still I can see that these two systems share many crucial formal feature. Willer ([100])'s non-monotonic dynamic account

I will first briefly summarize Veltman's Update Semantics, and then introduce my version of dynamic semantics using structure information states to analyze the two different roles played by *ought*-statements in our normative discourse and reasoning.

1.6.1 Veltman's Update Semantics

Veltman([96]) originally proposes his update semantics for epistemic *might*. So it starts with the following minimal language.

Definition 1 (Languages, $L_0^{\mathcal{A}}$ and $L_1^{\mathcal{A}}$) Let A be a set of atomic sentences, $p, q, r, ...$ The language $L_0^{\mathcal{A}}$ has $\mathcal{A}$ as its non-logical vocabulary and the logical operators, such as $\neg$, $\vee$, and $\wedge$, and parentheses,) and (, as its basic logical vocabulary.

A string of symbols ϕ is a sentence of $L_0^{\mathcal{A}}$ iff

(i) for some atomic sentence $p \in \mathcal{A}$, $\phi = p$;

(ii) for some sentence ψ of $L_0^{\mathcal{A}}$, $\phi = \neg\psi$;

(iii) for some sentences ψ and χ of $L_0^{\mathcal{A}}$, $\phi = \psi \wedge \chi$;

(iv) for some sentences ψ and χ of $L_0^{\mathcal{A}}$, $\phi = \psi \vee \chi$.

A language $L_1^{\mathcal{A}}$ has $\mathcal{A}$ and its logical vocabulary and an additional unary operator $\Diamond_e$ for *might*.

A string of symbols ϕ is a sentence of $L_1^{\mathcal{A}}$ iff for some sentence ψ of $L_0^{\mathcal{A}}$ $\phi = \psi$ or $\phi = \diamond_e\psi$

for deontic *ought*s is a type of update semantics, which solves Chisholm's Paradox by validating Factual Detachment. Compared to other deontic update semantics, the version of dynamic semantics I propose has a clear construction as a dynamic expansion of ordering semantics. And, most characteristically, it is motivated by philosophical and linguistic considerations on the dual nature of normative language.

Definition 2 (The Space of Possible Worlds, W) Let W is the set of all possible worlds, w. Each possible world w is an atomic valuation, a function from every atomic sentence in A to one truth value (1,0).

Definition 3 (Truth Sets/Propositions) For $p \in \mathcal{A}$, $[\![p]\!]^M = \{w \in W | w(p) = 1\}$. $[\![p]\!]$ is a set of possible worlds in which p is true. (M will be occasionally omitted for simplicity.)

Definition 4 (Information State) (i) σ is an *information state* iff $\sigma \subseteq W$; (ii) **0**, *the minimal state*, is the information state given by W; (iii) **1**, *the absurd state*, is the information state given by the empty set.[17]

For any sentence ϕ, $\mathbf{1}[\phi] = \mathbf{1}$. That means, in the absurd state every sentence is accepted, but no sentence is acceptable.

Definition 5 (Update System) An update system is a triple $\langle L, \Sigma, [*] \rangle$, where L is a language, Σ is a set of information states, σ, each of which is a subset of W, and $[*]$ is an update function. For any M, an update function $[*]$ is a function from a sentence of L to functions from information states to information states.

Definition 6 (Update Function) For every sentence ϕ of $L_1^{\mathcal{A}}$ and information state σ, $\sigma[\phi]$ is determined as follows:

$\sigma[p] = \sigma \cap [\![p]\!]$ $(p \in \mathcal{A})$

$\sigma[\neg\phi] = \sigma - \sigma[\phi]$

[17]Note that given the definition of W different sets of atomic sentences give rise to different sets of possible information states. So it would be more accurate to say A-information states and the A-minimal state.

$$\sigma[\phi \wedge \psi] = \sigma[\phi][\psi]$$

$$\sigma[\phi \vee \psi] = \sigma[\phi] \cup \sigma[\psi]$$

$$\sigma[\diamond_e \phi] = \begin{cases} \sigma, \textit{ if } \sigma[\phi] \neq \mathbf{1}; \\ \mathbf{1}, \textit{ if } \sigma[\phi] = \mathbf{1}. \end{cases}$$

Here $\sigma[p]$ is the state which results from updating σ to take account of p. Updating the information state σ with p has the effect of excluding all $\neg p$-worlds from σ. Notice that the function of epistemic modal statement behaves like a test. Updating σ with $\diamond_e p$ returns σ, if there is at least one p-world in σ; otherwise, it is an absurd state, $\mathbf{1}$, given by the empty set. For example, "It might be raining outside ($\Diamond_e r$)" is accepted by Jill only when her information state does not rule out the possibility that it is raining outside. So acceptance of an epistemic modal claim does not change her information state. Instead, it plays a role like a test; so it is accepted only when the information state passes the test.

Another central idea of a dynamic framework is the notion of *acceptance*. For an information state σ and a sentence ϕ, ϕ is accepted in the state σ when $\sigma[\phi] = \sigma$. This means, the information conveyed by ϕ is already assumed by σ. And, let's call the information state in which ϕ is accepted *the fixed point* of ϕ.

Definition 7 (Acceptance)

$\sigma \Vdash \phi$ *iff* $\sigma[\phi] = \sigma$

Veltman notes that the notion of acceptance in the dynamic system is normative rather than descriptive in that what are accepted in a certain information state should be accepted by any agent who are in that information state. If $\sigma[\phi] = \mathbf{1}$,

then an agent in σ should not accept ϕ; If $\sigma[\phi] = \sigma$, then an agent in σ should accept ϕ.

Consequence in a dynamic system can be defined in terms of acceptance as follows.

Definition 8 (Dynamic Consequence)

$\phi_1, \phi_2, ..., \phi_n \vDash \psi$ *iff* $\forall \sigma \in \Sigma$ *: if* $\sigma[\phi_1][\phi_2]...[\phi_n] = \sigma$*, then* $\sigma[\psi] = \sigma$*.*

This dynamic definition of consequence means that $\phi_1, ..., \phi_n \vDash \psi$, just in case every state of information which is a fixed point of $[\phi_1]...[\phi_n]$ also supports ψ. So the dynamic notion of consequence is the preservation of acceptance and sensitive to the order of premises.[18]

1.6.2 Structured Information States and Two *Ought*s

In the basic update semantics, an information state, σ, is modeled by a set of possible worlds. To incorporate the deontic ranking over possible worlds, I propose to replace a simple information state with a structured information state modeled by an (partially) ordered set of possible worlds. To keep track of live possibilities accepted in a structured information state and the relevant ranking over them, I define a structured information state τ with a pair of a set of possible worlds, σ, and a set of relevant normative principles, N, that determines the partial ordering among the possible worlds: $\tau = \langle \sigma, N \rangle$.

Definition 9 (Structured Information State τ) τ is a relevant agent(s)'s structured information state defined by a pair $\langle \sigma, N \rangle$.

[18] [96].

(i) $\sigma \subseteq W$: σ is a set of relevant live possible worlds

(ii)$N \subseteq \phi \times \psi$: N is a set of ordered pairs of propositions.

Note that a set of normative principles are modeled by a set of ordered pairs of proposition in order to capture the way conditional principles generate partial orderings. For example, the conditional moral principle expressed by "you ought to apologize if you break your promise($\mathbf{O}(a/b)$)," when accepted, generates a partial preference ordering ranking a-worlds higher than $\neg a$-worlds only among b-worlds.

Definition 10 (A Set of Normative Principles, N) For $\tau = \langle \sigma, N \rangle$, a normative principle expressed by a conditional *ought*-statement like "it ought to (or must) be ψ if ϕ, or a corresponding imperative ($!(\psi/\phi)$) is accepted in τ iff $\langle [\![\phi]\!], [\![\psi]\!] \rangle \in N$; an unconditional normative principle like "it ought to be ϕ" is accepted in τ iff $\langle W, [\![\phi]\!] \rangle \in N$.

For $\tau = \langle \sigma, N \rangle$, all possible worlds $u, v \in \sigma$ are related to each other in a way that either (i) u is strictly better than v ($u \prec_N v$), (ii) u and v are incomparable ($u \parallel_N v$), or (iii) u and v are competing ($u \sim_N v$). For $\tau = \langle \sigma, N \rangle$, N orders the possible worlds in σ as follows.

Definition 11 (Ordering by N)

For $\tau = \langle \sigma, N \rangle$, where $N = \{\langle [\![\phi_1]\!], [\![\psi_2]\!] \rangle ... \langle [\![\phi_n]\!], [\![\psi_n]\!] \rangle\}$, (for n, m, k = 1, 2, 3, ...)

(i) For $u, v \in \sigma$, $u \prec_N v$ iff there is at least one $\langle [\![\phi_m]\!], [\![\psi_m]\!] \rangle$ such that $u, v \in [\![\phi_m]\!], u \in [\![\psi_m]\!]$, and $v \notin [\![\psi_m]\!]$, but there is no $\langle [\![\phi_k]\!], [\![\psi_k]\!] \rangle$ such that $u, v \in [\![\phi_k]\!], v \in [\![\psi_k]\!]$ and $u \notin [\![\psi_k]\!]$.

(ii) For $u, v \in \sigma$, $u \parallel_N v$ iff there is no $\langle [\![\phi_m]\!], [\![\psi_m]\!] \rangle$ such that $u, v \in [\![\phi_m]\!], u \in [\![\psi_m]\!]$ and $v \notin [\![\psi_m]\!]$.

(iii) For $u, v \in \sigma$, $u \sim_N v$ iff there is $\langle [\![\phi_m]\!], [\![\psi_m]\!] \rangle$ such that $u, v \in [\![\phi_m]\!], u \in [\![\psi_m]\!]$ and $v \notin [\![\psi_m]\!]$, and there is $\langle [\![\phi_k]\!], [\![\psi_k]\!] \rangle$ such that $u, v \in [\![\phi_k]\!], v \in [\![\psi_k]\!]$ and $u \notin [\![\psi_k]\!]$.

Like the two contextual parameters f and g in Kratzer semantics, a pair $\langle \sigma, N \rangle$ shapes a partially ordered set of possible worlds. But the role played by this pair in my update semantics is very different from what $\langle f, g \rangle$ does in ordering semantics. Most importantly, $\langle \sigma, N \rangle$ represents a relevant information state, and the meaning of a sentence is the way it changes information states. By contributing to an information state $\langle \sigma, N \rangle$, utterances and uses of a sentence can change which possibilities are relevant in the context *and* the preference ranking over them. Also, and very importantly: by explicitly incorporating the set of relevant normative principles as another definer of information states, N, in a dynamic framework, we can capture two different mechanisms by which an information state can support an *ought*-statement. I'll now say a bit more about these two mechanisms.

I have argued that there are two kinds of *ought*-statements: axiological and deontological. My motivating observation is that although there is only one orthographic type *ought* in English, it can have different meanings and serve different functions in normative reasoning, depending on what kind of normative reason is intended to be conveyed: judgments about what is good or judgments about what is required. What a satisfactory semantic account of normative language should do is provide an adequate semantic account for both meanings of *ought*-statements after its possible ambiguity is resolved by the context of use.

To be clear, I do not claim that in ordinary English the context of use always decides between axiological and deontological readings of a particular use of *ought*; sometimes, the matter is simply unclear. Indeed, on my view, this is precisely why

we can read the Chisholm set in several different ways. Still, in most ordinary discourse, whether a certain use of an *ought*-statement is axiological or deontological is settled by the context of use, including the speaker's intentions. So exactly how we catch this intended meaning in everyday conversation remains as a pragmatic matter.

Although Kratzer semantics can in a sense represent the two meanings of *ought*-statements, invoking context to disambiguate between them, the two dimensions in structured information states afford us a better tool. That is because, as argued above, they capture two importantly different ways that a structured information state can support an *ought*-statement:

$$\tau \Vdash \mathbf{O}(\psi/\phi) \text{ iff for } \tau = \langle \sigma, N \rangle \text{ either } \sigma[\mathbf{O}(\psi/\phi)] = \sigma \text{ or } \langle [\![\psi]\!], [\![\phi]\!] \rangle \in N.$$

This means, one in a certain information state can accept or support an *ought*-statement either because it correctly describes what is the most desirable possibility in a relevant sense in the information state *or* because there is a corresponding normative principle accepted by the information state. This distinction nicely aligns with the different functions of axiological and deontological *ought*-statements: an axiological *ought*-statement describes what is the most desirable thing to do in a given circumstance, while a deontological *ought*-statement tells what is required by a relevant normative principle. Therefore, we can naturally draw the axiological and deontological *ought* distinction in terms of their different ways of updating a structured information state:

$$\tau[\mathbf{O}(\psi/\phi)] = \begin{cases} \langle \sigma[\mathbf{O}(\psi/\phi)], N \rangle, \text{if } \mathbf{O}(\psi/\phi) \text{ is an axiological } \textit{ought}\text{-statement;} \\ \langle \sigma, N \cup \langle [\![\psi]\!], [\![\phi]\!] \rangle \rangle, \text{if } \mathbf{O}(\psi/\phi) \text{ is a deontological } \textit{ought}\text{-statement.} \end{cases}$$

To keep track of different semantic functions that *ought*-statements play, let's introduce different notations for axiological and deontological *ought*-statements in this dynamic system: $\mathbf{O}_d(\psi/\phi)$ for deontological *ought*s and $\mathbf{O}_a(\psi/\phi)$ for axiological *ought*s.

Definition 12 (Deontic Language, $L_2^{\mathcal{A}}$) A string of symbols ϕ is a sentence of $L_2^{\mathcal{A}}$ iff (i) for some sentence ψ of $L_0^{\mathcal{A}}$, $\phi = \psi$; (ii) for some sentences ψ and χ of $L_0^{\mathcal{A}}$, $\phi = \mathbf{O}_a(\psi/\chi)$ or $\phi = \mathbf{O}_d(\psi/\chi)$.[19]

Definition 13 (Update Function for Structured Information States) For every sentence ϕ of $L_2^{\mathcal{A}}$ and information state τ, $\tau[\phi]$ is determined as follows.

$$\tau[p] = \langle \sigma[p], N \rangle = \langle \sigma \cap [\![p]\!], N \rangle$$

$$\tau[\neg\phi] = \langle \sigma[\neg\phi], N \rangle = \langle \sigma - \sigma[\phi], N \rangle$$

$$\tau[\phi \wedge \psi] = \langle \sigma[\phi \wedge \psi], N \rangle = \langle \sigma[\phi][\psi], N \rangle$$

$$\tau[\phi \vee \psi] = \langle \sigma[\phi \vee \psi], N \rangle = \langle \sigma[\phi] \cup \sigma[\psi], N \rangle$$

$$\tau[\boldsymbol{O}_a(\psi/\phi)] = \langle \sigma[\boldsymbol{O}_a(\psi/\phi)], N \rangle$$

$$\tau[\boldsymbol{O}_d(\psi/\phi)] = \langle \sigma, N[\boldsymbol{O}_d(\psi/\phi)] \rangle$$

Note that $\mathbf{O}_a(\psi/\phi)$ has an effect on σ and $\mathbf{O}_d(\psi/\phi)$ on N.

Updating a structured information state with a deontological statement amounts to adding $\langle [\![\phi]\!], [\![\psi]\!] \rangle$ to N. N as a set of ordered pairs determines the preference ranking over the possible worlds in σ.

[19] To set aside complex issues regarding the negations of modal sentences and compound sentence in this paper, I define the deontic language $L_2^{\mathcal{A}}$ with $L_0^{\mathcal{A}}$.

Definition 14 (Deontological *Ought*s : $\mathbf{O}_d(\psi/\phi)$)

$N[\boldsymbol{O}_d(\psi/\phi)] = N \cup \{\langle [\![\phi]\!], [\![\psi]\!] \rangle\}$

Updating the set of norms N with $\mathbf{O}_d(\psi/\phi)$ has the effect of adding a preference ordering. When $\tau[\mathbf{O}_d(\psi/\phi)] = \tau$, this information structure, τ, ranks $\phi\&\psi$-worlds over $\phi\&\neg\psi$-worlds among the accessible possible worlds in σ.

Axiological statements describe which a set of possible worlds are better than others, since they are about comparative betterness of possible worlds. So they behave similarly to conditionals or epistemic possibility modal statements in a dynamic semantics. Updating a structured information state with them amounts to testing whether the structured information is compatible with them.

Definition 15 (Axiological *Ought*'s : $\mathbf{O}_a(\psi/\phi)$)

$\tau[\boldsymbol{O}_a(\psi/\phi)] = \langle \sigma[\boldsymbol{O}_a(\psi/\phi)], N \rangle$*;*

$$\langle \sigma[\boldsymbol{O}_a(\psi/\phi)], N \rangle = \begin{cases} \langle \sigma, N \rangle, \textit{ if } \{w : w \in \sigma[\phi] \wedge \neg\exists v \in \sigma[\phi] : v \prec_N w\} \subseteq [\![\psi]\!]; \\ \emptyset, \textit{ otherwise.} \end{cases}$$

When an information state, τ supports ϕ, it can be presented as if ϕ has a static value relative to τ for the sake of convenience. We can define static values of both *ought*'s as follows.

When $\tau[\mathbf{O}_a(\psi/\phi)] = \tau$, $[\![\mathbf{O}_a(\psi/\phi)]\!]^\tau = 1$ when $\sigma[\phi] \neq \emptyset$.

When $\tau[\mathbf{O}_d(\psi/\phi)] = \tau$, $[\![\mathbf{O}_d(\psi/\phi)]\!]^\tau = 1$ when $\langle [\![\psi]\!], [\![\phi]\!] \rangle \in N$.

This static notation is useful when marking which subset of σ supports ϕ. For example, axiological *ought*-statements describe either the preference ordering of

possible worlds in σ or the subset of σ depending on a contextually salient domain. Updating τ with $\mathbf{O}_a(\psi/\phi)$ has the same effect to that of updating $\tau[\phi]$ with $\mathbf{O}_a(\psi/\top)$.

For ϕ, ψ in $L_0^{\mathcal{A}}$, $\langle \sigma, N\rangle[\mathbf{O}_a(\psi/\phi)] = \langle \sigma[\phi], N\rangle[\mathbf{O}_a(\psi/\top)]$ when $\sigma[\phi] \neq \emptyset$

It can be expressed by the identity of the static values.

$[\![\mathbf{O}_a(\psi/\phi)]\!]^{\tau} = [\![\mathbf{O}_a(\psi/\top)]\!]^{\tau[\phi]}$, when $\tau = \langle \sigma, N\rangle$.

Fixed Point (Acceptance), Consequence, and Consistency for structured information states, τ, are defined similarly as follows.

Definition 16 (Fixed Point for Structured Information States)
Information structure τ is a fixed point of ϕ iff $\tau[\phi] = \tau$

Definition 17 (Dynamic Consequence)
$\phi_1, \phi_2, ..., \phi_n \models \psi$ *iff* $\forall \tau \in T$ *: if* $\tau[\phi_1][\phi_2]...[\phi_n] = \tau$*, then* $\tau[\psi] = \tau$.

Given this definition although $\mathbf{O}_a(\psi_1 \wedge \psi_2/\phi)$ entails $\mathbf{O}_a(\psi_1\psi_2/\phi)$, $\mathbf{O}_d(\psi_1 \wedge \psi_2/\phi)$ does not entail $\mathbf{O}_d(\psi_1/\phi)$. The ordinary entailment relation does not hold for normative principles themselves, but for axiological *ought*-statements. We cannot generate a new rule from an existing set of rules, but infer what ought to be done.

Definition 18 (Consistency)
A sequence of sentences, $\phi_1, ..., \phi_n$ is consistent iff $\exists \tau : \tau[\phi_1]...[\phi_n] \neq \emptyset$

Note that given this definition of consistency $\mathbf{O}_d(p/\top)$ and $\mathbf{O}_d(\neg p/\top)$ are consistent while $\mathbf{O}_a(p/\top)$ and $\mathbf{O}_a(\neg p/\top)$ are not. This is not an odd result considering that our normative system can have conflicting enjoinments in practice, while the following conjunction of evaluative judgments cannot be made consistently: to realize p is the most desirable thing to do, and to realize $\neg p$ is the most desirable thing to do.

In sum, to say that the structured information state τ is a fixed point of $\mathbf{O}_a(q/p)$ means that it is adequate to state an axiological claim that it ought to be that q if p is adequate or felicitous; in other words, when one is in the information state $\tau[\mathbf{O}_a(q/p)] = \tau$, one believes that $p\&q$ is better than $p\&\neg q$. On the other hand, to say that the structured information state τ is a fixed point of $\mathbf{O}_d(q/p)$ means that it is adequate to state there is a norm that it ought to be that q if p; in other words, when one is in the information state $\tau[\mathbf{O}_d(q/p)] = \tau$, one prefers $p\&q$ over $p\&\neg q$, or accepts a norm preferring $p\&q$ over $p\&\neg q$.

Here are three ideas of the dynamic account sharing with ordering semantics; (i) an information state in this dynamic account is structured in that it is represented by a pair of a set of accessible possible worlds, σ, which is parallel to $\cap f(i)$ in Kratzer semantics, and a set of norms, N, which is parallel to Kratzer's ordering set, $g(i)$; (ii) updating deontological *ought*-statements to a structured information state, $\langle \sigma, N \rangle$, brings about a change of N, while axiological statements behave like a test on τ; and (iii) the static values of axiological statements can express the context sensitivity of modal statements depending on which subset of σ is under consideration.

1.6.3 Diagrams

The general idea of this approach can be easily illustrated with diagrams. In a diagram system, a solid line box represent a set of relevant possible worlds, σ, given a current information state; a dashed line box stands for the default information state relevant to a type of modality, and a circle stands for a non-empty set of possible worlds that satisfy the proposition written in the circle. Suppose that a current information state is the default information state with zero knowledge and preference. Such an information state does not rule out any possibilities and is indifferent among them. Updating σ with p has the effect of shrinking the solid line box to a set of relevant possibilities to the one compatible with p.

Suppose a structured information state τ_0 for which $\sigma = W$ and $N = \emptyset$; that is, the information state τ_0 is a set of all possible worlds in W, and there is no salient ranking over them. Such an information state τ_0 is illustrated as in Fig.1.1.

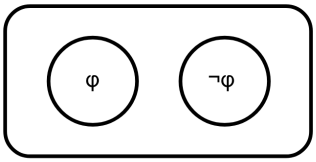

Figure 1.1: $\tau_0 = \langle \sigma, N \rangle$

τ_0 is not a fixed point of any axiological *ought*-statements since there is no relevant evaluative preference over possibilities.

Suppose that τ_0 is updated with $\mathbf{O}_d(\phi/\top)$. $\mathbf{O}_d(\phi/\top)$ is added to N, and $N[\mathbf{O}_d(\phi/\top)]$ ranks possible worlds in σ by connecting ϕ-worlds and $\neg\phi$-worlds and putting all ϕ-worlds on the left. Let's call this resulting information state τ_1 illustrated in Fig.1.2. $\tau_1 = \langle \sigma, N_1 \rangle$, where $N_1 = \{\langle W, [\![\phi]\!] \rangle\}$. τ_1 is the fixed point of $\mathbf{O}_a(\phi/\top)$; $\tau_1[\mathbf{O}_a(\phi/\top)] = \tau_1$, and thus, $[\![\mathbf{O}_a(p/\top)]\!]^{\tau_1} = 1$

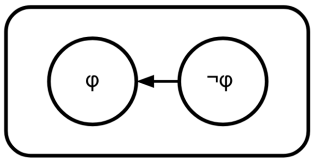

Figure 1.2: $\tau_1 = \tau_0[\mathbf{O}_d(\phi/\top)]$

When ϕ is not a tautology, adding $\mathbf{O}_d(\psi/\phi)$ to τ_0 has the effect of connecting the $\phi \wedge \psi$-worlds and $\phi \wedge \neg\psi$-worlds, putting the $\phi \wedge \psi$-worlds on the left, and leaving the rest of the worlds unconnected in σ as illustrated in Fig.1.3.

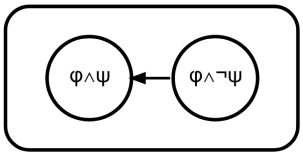

Figure 1.3: $\tau_2 = \tau_0[\mathbf{O}_d(\psi/\phi)]$

$\tau_2 = \langle \sigma, N_2 \rangle$ where $N_2 = \{\langle [\![\phi]\!], [\![\psi]\!] \rangle\}$. Thus, τ_2 is a fixed point of $\mathbf{O}_a(\psi/\phi)$; $[\![\mathbf{O}_a(\psi/\phi)]\!]^{\tau_2} = 1$. Note that adding a conditional deontological *ought*-statement to a set of norms induces a partial connection. It only takes effect on ϕ-worlds, and ranks ψ-worlds higher than $\neg\psi$-worlds among them. In this process, $\neg\phi \wedge \psi$-worlds and $\neg\phi \wedge \neg\psi$-worlds are left unconnected; it means that $\neg\phi \wedge \psi$-worlds and $\neg\phi \wedge \neg\psi$-worlds are incomparable with other worlds and with each other in σ.

Now let's see how we can take factual statements into the picture. Suppose $\tau_3 = \langle \sigma, N_3 \rangle$ where $N_3 = \{\langle W, [\![\phi]\!] \rangle, \langle [\![\neg\phi]\!], [\![\psi]\!] \rangle\}$, and update τ_3 with $\neg\phi$.

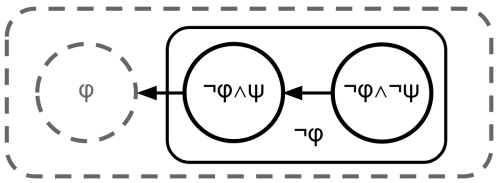

Figure 1.4: $\tau_3[\neg\phi] = \tau_0[\mathbf{O}_d(\phi/\top)][\mathbf{O}_d(\psi/\neg\phi)][\neg\phi] = \tau_1[\mathbf{O}_d(\psi/\neg\phi)][\neg\phi]$

Updating τ_3 with $\neg\phi$ has the effect of putting a solid line box around $\neg\phi$-worlds. $\tau_3[\phi]$ is the fixed point of $\mathbf{O}_a(\psi/\top)$ since ϕ-worlds are ruled out as adding $\neg\phi$. For

such τ_3, $[\![\mathbf{O}_a(\phi/\top)]\!]^{\tau_3} = [\![\mathbf{O}_a(\psi/\top)]\!]^{\tau_3[\neg\phi]} = 1$. So an axiological *ought*-statement describes the preference ranking among the possible worlds inside a solid line box. If there were no factual statements to consider, what one actually ought to do would have been identical to what ideally ought to do. We can distinguish what one ideally ought to do from what one actually ought to do with factual statement updates. The leftmost circle in the diagram is the set of ideal worlds without factual restriction; while the leftmost circle in a solid line box is the set of best worlds given the truth of $\neg\phi$. These diagrams illustrate how deontological statements order possible worlds, and what kind of axiological statements hold in a given domain. With this apparatus, we can distinguish what is required to be done by a set of relevant normative principles from what is the best thing can be done given what holds in a certain circumstance.

1.6.4 The CTD Trilemma and The Dynamic Solution

Given the two operators for *ought*, there are many different ways of reading the Chisholm set and the Extended Chisholm set. As we have seen, the key difference of the CTD Situation from the ATD and Overriding Situations is that (c1-3) are read as a deontological *ought*-statement.

(c1-d) $\mathbf{O}_d(h/\top)$

(c2-d) $\mathbf{O}_d(t/h)$

(c3-d) $\mathbf{O}_d(\neg t/\neg h)$

(c4) $\neg h$

(c5-a) $\mathbf{O}_a(\neg t/\top)$

This set depicts the structured information state like $\tau_5 = \langle \sigma, N_5 \rangle[\neg h]$, where $N_5 = \{\langle W, [\![h]\!] \rangle, \langle [\![h]\!], [\![t]\!] \rangle, \langle [\![\neg h]\!], [\![\neg t]\!] \rangle\}$, which formally represents the set of deontological statements: $\mathbf{O}_d(h/\top)$, $\mathbf{O}_d(t/h)$, $\mathbf{O}_d(\neg t/\neg h)$. τ_5 is illustrated as follows.

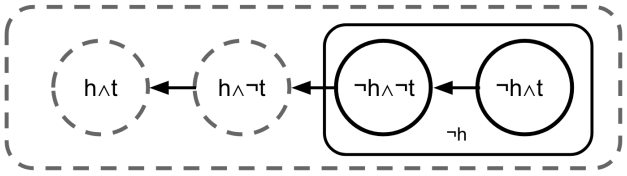

Figure 1.5: $\tau_5 = \tau_0[\mathbf{O}_d(h/\top)][\mathbf{O}_d(t/h)][\mathbf{O}_d(\neg t/\neg h)][\neg h]$

The possible worlds in τ_5 are completely connected by N_5. The factual statement $\neg h$ narrows down a practically relevant set into a set of $\neg h$-worlds – -the solid line box in Fig.1.5.[20]

The structured information state τ_5 adequately represents the Violation situation of the Chisholm set in that it supports (is a fixed point of) the above version of Extended Chisholm set, (c1-d), (c2-d), (c3-d), (c4) and (c5-a).

$$\tau_5[\mathbf{O}_d(h/\top)] = \tau_5$$
$$\tau_5[\mathbf{O}_d(t/h)] = \tau_5$$
$$\tau_5[\mathbf{O}_d(\neg t/\neg h)] = \tau_5$$
$$\tau_5[\neg h] = \tau_5$$
$$\tau_5[\mathbf{O}_a(\neg t/\top)] = \tau_5$$

This result meets all three desiderata of the CTD Trilemma. Jones' violation of his duty to go to help is expressed by the conjunction, $\mathbf{O}_d(h/\top) \wedge \neg h$, and τ_5 is a fixed point of both $\mathbf{O}_d(h/\top)$ and $\neg h$. Second, τ_5 is the fixed point of (c5-a) $\mathbf{O}_a(\neg t/\top)$. Finally, this extended Chisholm set is free from any tensions; τ_5 is

[20]For τ_5, the following axiological statements hold: $[\![\mathbf{O}_a(h/\top)]\!]^{\tau_5} = [\![\mathbf{O}_a(h \wedge t)]\!]^{\tau_5} = [\![\mathbf{O}_a(t/h)]\!]^{\tau_5} = [\![\mathbf{O}_a(\neg t/\neg h)]\!]^{\tau_5} = [\![\mathbf{O}_a(\neg t/\top)]\!]^{\tau_5[\neg h]} = 1$.

the fixed point of all five statements: (c1-d), (c2-d), (c3-d), (c4) and (c5-a). Also, the existence of τ_5 proves that these five sentences are compatible. This dynamic account draw the wanted distinction between axiological and deontological *ought*-statements effectively and meet the CTD Trilemma.

Moreover, this dynamic system can explain why reading (c1)-(c3) axiologically generates the CTD Trilemma. Suppose that there is a structured information state $\tau_6(i)$ which supports the axiological reading of the Chisholm set: (c1-a), (c2-a) and (c3-a). So $\tau_6(i)$ is a fixed point of $\mathbf{O}_a(h/\top), \mathbf{O}_a(t/h)$ and $\mathbf{O}_a(\neg t/\neg h)$. Once $\tau_6(i)$ is updated with (c4) $\neg h$, $\tau_6(i)[\neg h]$ stops supporting (c1-a)$\mathbf{O}_a(h/\top)$, but supports (c5-a) $\mathbf{O}_a(\neg t/\top)$, because the update shrinks the relevant set of possible worlds to a set of $\neg h$-worlds. There is no non-empty information states that support the following axiological version of the Extended Chisholm set.

(c1-a) $\mathbf{O}_a(h/\top)$

(c2-a) $\mathbf{O}_a(t/h)$

(c3-a) $\mathbf{O}_a(\neg t/\neg h)$

(c4) $\neg h$

(c5-a) $\mathbf{O}_a(\neg t/\top)$

Suppose that τ_6 is the fixed point of (c1-a), (c2-a) and (c3-a). After (c4) is accepted by τ_6, this information state stops supporting (c1-a), like in the Overriding Situation. (c1-a) is overridden by (c4). On the other hand, before accepting (c4) in this information state, τ_6 supports (c1-a), (c2-a),(c3-a) and (c6-a) like the ATD Situation. In this way, the dynamic semantic account provides the consistent explanation of other situations as well.

1.7 Conclusion

So far we have critically discussed ordering semantics for deontic *ought*s with the CTD Trilemma. This new puzzle reveals the limitations of ordering semantics, and highlights the two functions served by *ought*s in our normative discourse and reasoning. To capture them I have proposed a version of dynamic semantics free from the CTD Trilemma. The shift from a static system to a dynamic one might be considered a big leap. However, the proposed version of dynamic semantics opens up the possibility of dynamic implementation of ordering semantics by preserving static meanings of modal claims. Moreover, the significance of dynamic approach deserves to be evaluated in the larger context of linguistic data, moral theories and metaethics. It is widely recognized that the normative use of language plays two different roles in our normative discourse and reasoning: the prescriptive and descriptive features. Dynamic semantics is an ideal semantic framework for formally representing this dual nature of normative language as well. The dynamic understanding of meaning and communication has great advantages in capturing various insights into our natural language use and semantically theorizing complex reasoning and inferences.

CHAPTER 2

MORAL DILEMMAS IN ORDERING SEMANTICS

The possibility of moral dilemmas is one of many classic challenges to a formal account of normative language. In this paper, I show how the standard semantics for *ought*-statements, ordering semantics, can, in fact, meet the challenges raised by the possibility of moral dilemmas and reasoning with conflicting *ought*-statements. After examining several possible modifications of ordering semantics, I conclude that to give a satisfactory representation of moral dilemma situations and of our reasoning with conflicting obligations, a formal account needs to be equipped with two kinds of *ought*-operators.

First, I introduce the notion of moral dilemmas and briefly explain why the possibility of moral dilemmas causes trouble for moral theories and formal accounts of normative language. In sec 3., I discuss how a moral dilemma can be understood in ordering semantics. Ordering semantics implements the distinction between what is believed and what is desired (or preferred) using two contextual parameters. This apparatus opens up the possibility of representing a moral dilemma without incurring deontic explosion. The simple account using these contextual parameters locates a conflict between moral principles in the right place, but it does not provide satisfactory predictions about moral dilemma situations. In sec 4., I introduce and discuss two alternative ways of capturing moral dilemmas in the ordering semantic framework: the conflict and disjunctive accounts inspired by Horty ([48]). I show that each account gets some aspects of our moral reasoning with conflicting obligations right. But each has its own limitations. The strengths and weaknesses of each account reveal that there are in fact two distinct types of *ought*-statements. Finally, I explore the possibility of a third account that combines both the conflict

and disjunctive accounts of moral dilemmas and conclude that the advantages of this unified view adequately motivate us to introduce two kinds of *oughts* in ordering semantics. Developing a picture that fully integrates the two kinds of *ought*-statements is a project for further research. Still this paper aims to show that in order to capture our reasoning with conflicting duties and obligations in a formal system, we need to acknowledge two *oughts* in normative reasoning and introduce two *oughts* into a semantic framework. This paper shows how it can be done in ordering semantics.

2.1 Introduction: Moral Dilemmas

Most of us have experienced having to make a hard choice between more than one alternative, and we call some subset of such situations "moral dilemmas." An agent in a moral dilemma situation is required to perform two (or more than two) actions but cannot do both, and whatever choice she makes in this situation will violate one of her duties or obligations. So a dilemma situation for an agent S can be represented by a pair of conflicting *ought*s for S: "S ought to do A" and "S ought not to do A"; or "S ought to do A" and "S ought to do B" when S cannot do both A and B.

The possibility of genuine moral dilemmas poses a significant problem not only for moral theories but also for formal accounts of normative language. For ethicists, the possibility of genuine moral dilemmas means that the moral theory that they endorse fails to help us figure out what actions are right and wrong in all situations. For logicians, semanticists, and philosophers of language, who care about formal accounts of natural language and are out to model human reasoning using such accounts, accepting the possibility of genuine moral dilemmas means

that a formal representation of normative language should be able to tolerate a sort of inconsistency or at least an incompatible set of *ought*-statements, without unacceptable consequences like *deontic explosion*—allowing absolutely everything to count as obligatory.[1]

Of course, at the end of the day it could turn out that there are no genuine moral dilemmas. It could be that all of the putative moral dilemmas that we experience are due to the failure of seeing the priority or hierarchy in different duties and obligations, or the failure of weighing the importance of the different roles and responsibilities that an agent adopts. In this case, an apparent dilemma situation is simply a situation in which we do not know which of various competing principles is more important or supposed to weigh more. In that case, an adequate formal account for apparent moral dilemmas should be about our reasoning under uncertainty. It does not have to be an inconsistency-tolerant system. Still, this formal account should be able to accommodate the possibility of believing or accepting two seemingly conflicting moral principles to an equal degree.

Moral theoretical worries and our epistemic limitations as human beings do not reduce the importance of developing a formal account of normative language and reasoning that can represent moral conflicts and our reasoning in dilemma situations. All that these considerations rather show is that whether a true moral dilemma really exists or not and whatever its nature may be (ontological or epistemic), there should be some formal way of capturing our reasoning about conflicting obligations in what appears to us as a dilemma situation. Therefore, in this paper, I assume that moral dilemmas are genuinely possible, and understand a dilemma situation as one in which at least two inconsistent *ought*-statements

[1] According to the principle of explosion in classical logic, any statement can be proven from a contradiction. In a deontic logic, the deontic version of this principle holds: from any contradictory *ought*-statements anything can be proven to be obligatory.

hold, or at least, in which we have strong reasons to regard both of them as true and as telling us what to do.

Before discussing some formal difficulties that the possibility of moral dilemmas poses, let me flag some characteristic features of moral dilemmas that distinguish them from other types of hard choices. An agent in a dilemma situation faces a hard choice between equally undesirable or unacceptable options because there will inevitably be unwanted violation of a moral principle that the agent endorses. Suppose that you have decided to put Saturday afternoons to good use by helping others. You can volunteer at the community center or at a local hospital in your community. How should you spend your Saturdays? When both activities are equally important and hard to compare, it is a hard choice. However, this is not a dilemma; it is merely a case of overabundant choices. In a dilemma situation, an agent faces a choice of two options with equally undesirable and unpalatable consequences, and so she has good reasons not to want to take either. One's choice in a true dilemma situation is not only undesirable but also leaves *moral residue*—that is, "moral distress" that is "long-lasting and powerfully integrated into one's thoughts and views of the self [and] can be damaging to the self and one's career."[2] Moral residue that an agent would have after making a decision in a dilemma situation is not simply due to the negative consequences of a choice. This means that not all choices in which both options have serious problems are moral dilemmas. Suppose that two emergency patients have just arrived in ER, and the only available doctor there can start an operation on only one patient. Either choice would result in another patient's death. Given that there is no reason to pick one patient over the other, both choices equally have a bad result: the death of the other patient. In this case, flipping a coin might be the only solution. Importantly,

[2]Epstein and Delgado ([22])

however, neither choice of the doctor violates her duties or obligations as a doctor. As a moral agent, she might feel some moral compunction about not being able to save both, but there is no reason for her to regret and feel guilty about making a choice—in a sense, either choice will be morally desirable despite the fact that each entails bad consequences. On the other hand, in a moral dilemma situation, each choice is morally undesirable in that by choosing either one of the options, the agent violates one of her duties or obligations. In sum, an agent faces a moral dilemma if she has to make a choice between two incompatible alternatives, but either choice would violate some of her obligations.

In this paper, I will argue that the current standard semantics of *ought*-statements, namely, the ordering semantics proposed and developed by Angelika Kratzer and David Lewis,[3] does not adequately represent dilemma situations and our reasoning with conflicting *ought*-statements. In particular, I will show that there are two straightforward modifications that ordering semantics can take to make some room for the possibility of moral dilemmas: the *conflict* account and the *disjunctive* account. Each account captures some important aspects of dilemma situations, but each has characteristic drawbacks. These two approaches to dealing with conflicting *ought*-statements reveal that there are two kinds of *ought*-statements in normative discourse and reasoning. I will then explore the possibility of implementing both accounts in the standard semantics, assessing the costs and advantages of this third route.

Before I start discussing the limitations of the orthodox semantics, let me present the obvious formal difficulties posed by the possibility of moral dilem-

[3]Ordering semantics is the standard semantics for modals and conditionals today in linguistics and philosophy of language, and most famously developed by Kratzer([55], [56],[57]), Lewis([60], [61],[62]), Stalnaker([87]), von Wright([98], [40], [41]), Hansson([35]) and Føllesdal and Hilpinen([25]).

mas. Readers who are already familiar with the problems of dilemmas in Standard Deontic Logic may want to skip Section 2.

2.2 Moral Dilemmas in Standard Deontic Logic

Standard Deontic Logic (SDL) is a normal modal logic with a deontic necessity operator '**OB**' for obligations, and is axiomatized as follows:

A1	All tautologies	
A2	$\mathbf{OB}(\phi \supset \psi) \supset (\mathbf{OB}\phi \supset \mathbf{OB}\psi)$	(**OB**-K)
A3	$\mathbf{OB}\phi \supset \neg\mathbf{OB}\neg\phi$	(**OB**-Dual)
R1	If $\vdash \phi$ and $\vdash \phi \supset \psi$ then $\vdash \psi$.	(MP)
R2	If $\vdash \phi$ then $\vdash \mathbf{OB}\phi$.	(**OB**-Nec)

Since in SDL ordinary if-statements are expressed by material conditionals, there are two possible ways that a material conditional and a deontic necessity operator can combine to express conditional obligations: "$\mathbf{OB}(\phi \supset \psi)$" and "$\phi \supset \mathbf{OB}(\psi)$." As a normal propositional logic, all propositional calculi (PC) hold for negation, conjunction, disjunction and material conditional as in other propositional logics. Given this, it is easy to see how the first type of conflicting *ought*s, $\mathbf{OB}(\phi)$ and $\mathbf{OB}(\neg\phi)$, lead to contradiction:

1.	$\mathbf{OB}(\phi)$	
2.	$\mathbf{OB}(\neg\phi)$	
3.	$\neg\mathbf{OB}(\neg\phi)$	1 by (**OB**-Dual)
4.	$\perp$	2, 3.

The second type of dilemma involves a situation in which an agent *ought* to do A, *ought* to do B, but cannot do both A and B. If we have the dual principle

in modal logic and the following intuitive principle of deontic logic, then the conjunction of $\mathbf{OB}(\phi)$ and $\mathbf{OB}(\psi)$ leads to a contradiction when an agent cannot do both.

Dual

$\vdash \Box(\phi) \supset \neg\Diamond(\neg\phi)$

Principle of Deontic Logic (PD)

$\vdash \Box(\phi) \supset \mathbf{OB}(\phi)$

1.	$\mathbf{OB}(\phi)$	
2.	$\mathbf{OB}(\psi)$	
3.	$\neg\Diamond(\phi \wedge \psi)$	
4.	$\Box(\neg\phi \vee \neg\psi)$	3 by (Dual)
5.	$\Box(\phi \supset \neg\psi)$	Propositional Calculus (PC)
6.	$\mathbf{OB}(\phi \supset \neg\psi)$	5 by (PD)
7.	$\mathbf{OB}(\phi) \supset \mathbf{OB}(\neg\psi)$	6 by (**OB**-K)
8.	$\mathbf{OB}(\neg\psi)$	1, 7 by (MP)
9.	$\neg\mathbf{OB}(\neg\psi)$	2 by (**OB**-Dual)
10.	$\perp$	2,9.

When ϕ entails $\neg\psi$ or $\neg\psi$ entails ϕ, it is even simpler to show the inconsistency in the second type of dilemma.

1. $\mathbf{OB}(\phi)$
2. $\mathbf{OB}(\psi)$
3. $\phi \supset \neg\psi$
4. $\mathbf{OB}(\phi \supset \neg\psi)$ 3 by (**OB**-Nec)
5. $\mathbf{OB}(\phi) \supset \mathbf{OB}(\neg\psi)$ 4 by (**OB**-K)
6. $\mathbf{OB}(\neg\psi)$ 1, 5 by (MP)
7. $\neg\mathbf{OB}(\neg\psi)$ 2 by (**OB**-Dual)
8. $\bot$ 2,7.

Both of these proofs appeal to (**OB**-Nec). However, it is often controversial whether what is necessary is obligatory. For example, 2+2=4 is a necessary truth; but it is not even very intelligible to say that it is obligatory that 2+2=4. Here is another way of showing the inconsistency generated by a dilemma without appealing to (**OB**-Nec). In this proof, the "Ought-Implies-Can" Principle, a moral principle widely accepted by ethicists, plays a central role.

The "Ought-Implies-Can" Principle (OIC)

$\vdash \mathbf{OB}(\phi) \supset \mathbf{CAN}(\phi)$.

Consider the following type of a dilemma situation: an agent ought to do ϕ and ought to do ψ, but cannot do both ϕ and ψ; so "$\mathbf{OB}(\phi)$," "$\mathbf{OB}(\psi)$," and "$\neg\mathbf{CAN}(\phi \wedge \psi)$" hold.

1. $\mathbf{OB}(\phi)$
2. $\mathbf{OB}(\psi)$
3. $\neg\mathbf{CAN}(\phi \wedge \psi)$
4. $\mathbf{OB}(\phi \wedge \psi)$ 1, 2 by (AND: the Aggregation Principle)
5. $\mathbf{OB}(\phi \wedge \psi) \supset \mathbf{CAN}(\phi \wedge \psi)$ 4 by (OIC)
6. $\mathbf{CAN}(\phi \wedge \psi)$ 4,5 by (MP)
7. $\bot$ 3,6.

The logical problem of two conflicting *ought*-statments in a dilemma situation also can be understood in terms of deontic explosion (DEX), the deontic version of the Principle of Explosion.

Principle of Explosion (EXP)

$\vdash (\phi \wedge \neg\phi) \supset \psi$, for any ψ.

Principle of Deontic Explosion (DEX)

$\vdash (\mathbf{OB}(\phi) \wedge \mathbf{OB}(\neg\phi)) \supset \mathbf{OB}(\psi)$, for any ψ.

Goble ([32]) argues that to avoid (DEX) at least one of the following principles ought to be rejected or restricted: (RM) and (AND).

Rule of Monotonocity (RM) [4]

If $\vdash \phi \supset \psi$, then $\mathbf{OB}(\phi) \supset \mathbf{OB}(\psi)$.

The Aggregation Principle (AND)

$\vdash (\mathbf{OB}(\phi) \wedge \mathbf{OB}(\psi)) \supset \mathbf{OB}(\phi \wedge \psi)$.

[4]The Rule of Monotonocity (RM) follows from (**OB**-K) and (**OB**-Nec). Intuitively, (RM) just says that if doing ϕ brings about ψ, and if ϕ is obligatory (morally required), then ψ is obligatory (morally required).

1. $\mathbf{OB}(\phi)$
2. $\mathbf{OB}(\neg\phi)$
3. $\mathbf{OB}(\phi \wedge \neg\phi)$ 1,2, by (AND)
4. $(\phi \wedge \neg\phi) \supset \psi$ (EXP) for any B
5. $\mathbf{OB}(\phi \wedge \neg\phi) \supset \mathbf{OB}(\psi)$ 4 by (RM)

To develop deontic logic systems that accommodate the possibility of moral dilemmas without deontic explosion or inconsistency, different authors bring different principles under suspicion: (Dual), (OIC), (EXP), (RM) and (AND). Many have suggested that a dilemma-friendly logic cannot have either (RM) or (AND) in its full strength. Some even reject (DEX) itself.[5] For a systematic summary of different solutions to deontic dilemmas see Goble ([32]).

2.3 Moral Dilemmas in Ordering Semantics

We have seen that moral dilemmas expressed by a pair of conflicting *ought*-statements result in inconsistency or deontic explosion in SDL, and to accommodate moral dilemmas in a formal system, some of SDL's widely-accepted principles need to be rejected or restricted. In this section, I am going to review the formal difficulty posed by the possibility of moral dilemmas in ordering semantics, and the simple account that ordering semantics can provide to capture moral dilemma situations.

[5]Lemmon ([59]) rejects (OIC). Williams ([101]) and van Fraassen ([27]) accept (OIC), but require modifications of (AND) for deontic *ought*'s. Horty ([46], [47], [48]) argues that (AND) cannot be fully rejected in a formal system to accommodate deontic dilemmas. Jackson ([49]) Hansson ([36], [37]) and Goble ([30], [31]) call (RM) into question. Goble ([32]) proposes a modification of (RM) for moral reasoning to enable deontic logic to accommodate moral dilemmas without inconsistency or (DEX).

2.3.1 Conflicting Oughts in Ordering Semantics

Ordering semantics treats a moral (or normative) *ought*-statement as a deontic modal claim that describes what is the case in the possible worlds closest to the deontic ideality among relevant possible worlds. This idea is systematically rendered by two contextual parameters, a modal base and an ordering source. A modal base, $f(i)$, is a function mapping an index world i to a set of propositions. An ordering source, $g(i)$, is a function mapping i to a set of propositions that determines the ideals that are relevant to i. These two contextual parameters allow us to take two core ideas in modal reasoning into a semantic framework: what are the relevant possibilities in a given context and which possibilities are more preferred in that context. A deontic *ought*-statement, "it ought to be the case that ϕ" is represented by "$\mathbf{O}(\phi)$" with a propositional deontic modal operator "$\mathbf{O}$."[6] The meaning of deontic *ought*-statements in our normative discourse and reasoning is analyzed as follows in deontic ordering semantics:

For a language $\mathcal{L}$ and $\mathcal{M} = \langle \mathcal{W}, f, g, [\![\quad]\!] \rangle$,

$[\![\mathbf{O}(\phi)]\!]_i^{f,g} = 1$ iff $\{w : w \in \cap f(i) \wedge \neg \exists v \in \cap f(i) : v \prec_{g(i)} w\} \subseteq [\![\phi]\!]$,

where $\cap f(i) = \{w : w \in \cap f(i)\}$.

$\mathbf{O}(\phi)$ is true if and only if in the domain of accessible possible worlds there are some ϕ-worlds that are closer to the deontic ideality (or deontically better) than any $\neg\phi$-world.

[6] Given the way of using the propositional operator "$\mathbf{O}$" in ordering semantics, at the semantic level there is no meaningful difference between "It ought to be that if he does go he tells them he is coming" and "If he goes then he ought to tell them he is coming" . Also, in ordering semantics there is no distinction between agentive and non-agentive *ought*-statements. Some argue that this distinction corresponds to the grammatical distinction between "S ought to do/be x" and "It ought to be the case that p" (where S is an agent, ϕ is an action and p is a proposition): for example, "Sue ought to keep her promise" and "It ought to be that Sue keeps her promise." For further discussions on this distinction and possible modifications of ordering semantics based on this distinction, see Finlay and Snedegar ([24]) and Chrisman ([18]).

Consider, for example, the sentence "Jones ought to help his neighbors", or, more stiltedly, "It ought to be the case that Jones helps his neighbors." Let 'h' stand for "Jones helps his neighbors." Then in our formalism, $\mathbf{O}(h)$ is true just in case the possible worlds in which Jones helps his neighbors (h-worlds) are deontically better—closer to the contextually-set ideal world—than any possible world in which Jones does not ($\neg h$-worlds). As in SDL, in ordering semantics neither a pair $\mathbf{O}(\phi)$ and $\mathbf{O}(\neg\phi)$ nor a pair $\mathbf{O}(\phi)$ and $\mathbf{O}(\psi)$, where ϕ entails $\neg\psi$, can be true at the same time. However, ordering semantics *can* capture the conflicting obligations or moral principles that generate conflicting *ought*-statements. It does this by registering them in an ordering source for evaluating modal claims. Below, we will examine whether this way of accommodating conflicting obligations and normative principles in a formal account is satisfactory, and whether it lets us adequately represent our normative reasoning about moral dilemma situations.

2.3.2 Conflicting Obligations in Ordering Semantics

Kratzer, one of the major developers of ordering semantics along with David Lewis, points out that we routinely deal with inconsistent desires or conflicting commands. And, she identifies the phenomenon as a conflict between the relevant facts and what one wants to do.

> In our case, both conversational backgrounds assign consistent sets to the worlds compatible with our scenario. The facts are consistent. Facts always are. And what I want is consistent. too. Yet there is a conflict between the relevant facts and what I want. Not everything I want can be realized ([57]; 648).

She goes on to provide a diagnostic explanation of why the standard analysis fails to handle inconsistent desires and preferences as follows:

> The standard analysis cannot account for cases of this sort. Lacking the distinction between ordering source and modal base, it would have to lump together facts and desires. This would simply result in an inconsistent set. All necessity statements would come out true. All possibility statements would come out false. On the new analysis, the relevant facts form the modal base. What I want constitutes the ordering source ([57]; 648).

This inspiring observation opens up the possibility of formally presenting the inconsistent obligations that humans sometimes have. In the ordering semantics framework, we can model a dilemma situation by having the ordering source include two (conflicting) preference rankings and having the modal base exclude all possible worlds that satisfy both preference rankings. Let's call this approach of incorporating conflicting obligations in an ordering semantics a *simple account* for moral dilemmas. Suppose that I promised my daughter to drive her to her summer camp, and promised my colleague that I would substitute-teach his class on Thursday. And, suppose it turned out that my daughter needed to get to her summer camp during the class. Since I have promised both people, I ought to drop my daughter to her summer camp in time and I ought to teach class for my colleague on Thursday. However, I cannot do both. In this situation, I have two conflicting duties, stemming from the different roles that I take and the commitments that I made in the past. According to the simple account, my duties in this situation are captured by the relevant deontic ordering source ($g(i)$), and what I can or cannot do is captured by a modal base ($f(i)$).[7] The relevant ordering source in my situation

[7]Kratzer ([57]; 647-9).

is one that includes both "I drive my daughter to her summer camp on Thursday (d)" and "I teach class on Thursday (c)"; $g(i) = \{d, c\}$. The relevant modal base in the situation in which I cannot do both is: $f(i) = \{\neg(d \wedge c)\}$. This means there are no $d \wedge c$-worlds in the domain of accessible possible worlds—all d-worlds are $\neg c$-worlds and all c-worlds are $\neg d$-worlds. Although the simple approach captures one's obligations and duties in a relevant ordering source that contributes to the meaning of *ought*-statements in the context, the simple approach does not allow us the most straightforward way of describing my dilemma situation: (1)-(3).

(1)	I ought to drive my daughter to her summer camp on Thursday.	$\mathbf{O}(d)$
(2)	I ought to teach class for my colleague on Thursday.	$\mathbf{O}(c)$
(3)	I cannot do both.	$\neg\Diamond(d \wedge c)$

In the simple account, my predicament is not the situation where (1)-(3) hold. Given the ordering source and the modal base of my situation, the possible worlds in which I drive my daughter to her summer camp and the possible worlds in which I teach class for colleague on Thursday are *equally* close to the normative ideals for me; they are both as close as I can get. Therefore, neither (1) nor (2) holds, contrary to the intuitive characterization of my situation, (1)-(3).

2.3.3 Limitations of the Simple Account

It is an ingenious idea to represent duties and obligations separately with two contextual parameters: a modal base and ordering source. The simple account that directly incorporates this idea allows ordering semantics to avoid deontic explosion when representing a dilemma. However, the result of this simple account is odd because my dilemma situation cannot be described by the set of *ought*-statements

corresponding to my duties and obligations, such as "I ought to drive my daughter to her summer camp on Thursday ($\mathbf{O}(d)$)" and "I ought to teach class on Thursday ($\mathbf{O}(c)$)." It is particularly unsatisfactory given that ordering semantics is natural language semantics; *prima facie*, natural languages *do* expresses normative principles such as duties and obligations via *ought*-statements. Any ordinary speaker would not hesitate to summarize my situation using both (1) and (2). The simple account thus faces the worry that it inadequately represents situations of moral dilemma or normative conflict. Moreover, this simple account has another unnatural consequence: in my predicament, (5) and (6) hold, instead of (1) and (2).

> (5) It is not the case that I ought to drive my daughter to her summer camp on Thursday.
> (6) It is not the case that I ought to teach the class on Thursday.

Again, this is because on the simple account, it can only be that I ought to drive my daughter if every world in which this happens is closer to the ideal than every world in which it does not. Lassiter ([58]) expresses this type of worry as follows:

> The entire problem of moral conflict is that there are sometimes conflicting *ought*-statements which are simultaneously true; what Kratzer gives us instead is a semantics which renders them both false. (...) We wanted to find a logic for *ought* and should that makes it possible to model moral conflicts, but what we have here is one which simply ignores them ([58]; 151).

Ordering semantics can rightly tease apart two factors in our normative reasoning, factual and normative elements ("facts and desires" in her terminology), and

captures this distinction with the two contextual parameters f (modal base) and g (ordering source). However, simply assuming that the ordering source captures all of the relevant obligations does not allow us the most natural way of describing dilemma situations and our reasoning with conflicting *ought*-statements. This result is not only unsatisfactory from the standpoint of natural language semantics, but also conceptually problematic, especially for those who endorse the strong connection between *ought*s and reasons.[8] If we assume that *ought*-statements play motivating, reason-giving, and action-guiding roles in discourse and reasoning, then the situation characterized by (5) and (6) seems to be one in which I have normative reason neither to drive my daughter to her summer camp nor to teach class on Thursday. It is possible that I get paralyzed in the face of these conflicting duties and hard choices. But it is not the case that I am unmotivated or lack reasons for either action. On the contrary, in fact I have (equally strong) moral or normative reasons to do both.[9] After all, if I do nothing, then both my friend and my daughter will be hurt. Surely this is relevant to my motivations and my reasons for acting. The simple account accommodates the fact that normative principles and commands can conflict with each other and be inconsistent by incorporating them into an ordering source. However, we saw that this approach makes some unsatisfactory predictions about moral dilemma situations, which conflicts with the most straightforward way of describing one's dilemma. In the next section, I am going to discuss two possible modifications of ordering semantics to provide a better representation of moral dilemma situations.

[8]In fact, this connection is not free of controversy. While Broome ([12]) endorses the strong connection between oughts and reasons: the basic normative concept is that of ought-facts, and that reasons are to be explained by their role in explaining ought-facts. Raz ([81]) rejects this strong connection based on conceptual considerations. Bronfman and Dowell ([10]) and contextualists deny this connection between ought and normative reasons.

[9]The meaning of an *ought*-statement in this sense cannot be used for action-guiding purposes.

2.4 The Conflict and Disjunctive Accounts in Ordering Semantics

I am going to present two possible approaches that ordering semantics can take in order to formally represent dilemma situations in the ordering semantic framework: the disjunctive account and the conflict account. According to the disjunctive account, my predicament in the above example is the situation in which I *ought* either to drive my daughter to her summer camp or to teach class on Thursday. According to the conflict account, my predicament is the situation in which I ought to drive my daughter to her summer camp, and I ought to teach class on Thursday. These two possible accounts are inspired by Horty ([48])'s two deontic logics, which are intended to model two views in the debate over the possibility of conflicting all-things-considered *ought*s.[10]

To introduce and contrast these two possible modifications, let's take Sartre's Dilemma as an example of a moral dilemma.

2.4.1 Sartre's Dilemma

Consider the following story from Sartre's "Existentialism in a Humanism."

> One day student came to Sartre for advice about his predicament. He faced the decision between joining the Free French forces to fight for the liberation

[10]Horty ([48]) presents the conflict account and disjunctive account in his default logic system to compare two different strategies for reasoning with two conflicting *prima facie ought*s and to carefully examine the possibility of all-things-considered conflicts. Foot ([26]), Thomson ([92]), Pietroski ([71]) and Brink ([8]) endorse the view that there are moral conflicts among *prima facie ought*s, but not all-things-considered oughts.

of his country and remaining with his mother, who is completely dependent on him.

Not only is the predicament of Sartre's pupil, say Jean, possible, but also we easily understand the difficulty of handling such a situation because we all sometimes face a choice between two actions which is such that either decision will violate some of our duties and obligations. Jean faces a situation in which it is not possible for him to fulfill his two obligations. So the most straightforward way of representing his predicament would be to use the following set of sentences:

(s1)	Jean *ought* to join the Free French forces.	$\mathbf{O}(f)$
(s2)	Jean *ought* to remain with his mother.	$\mathbf{O}(r)$
(s3)	If Jean joins the forces, he does not remain with his mother, and if he remains with his mother, he does not join the forces.	$(r \rightarrow \neg f) \wedge (f \rightarrow \neg r)$

While the set of sentences (s1)-(s3) is the most straightforward way of describing the situation, in ordering semantics these three sentences cannot all be true at the same time. (s3) captures the fact that Jean cannot both join the Free French forces and remain with his mother. So to fulfill both obligations expressed by (s1) and (s2) is not practically possible for him in this situation. Despite this, our moral intuitions (and Jean's!) tell us that Jean's obligation to join the Free French forces and his obligation to remain with his mother both still hold. Yet in ordering semantics, when (s3) holds, both (s1) and (s2) cannot both be true. For $(r \rightarrow \neg f) \wedge (f \rightarrow \neg r)$ is true iff all accessible r-worlds are $\neg f$-worlds, and all accessible f-worlds are $\neg r$-worlds. This means that there are no accessible possible worlds in which both r and f hold. Therefore, there are no deontically ideal possible worlds among the accessible possible worlds in which both f and r

hold. Consider the truth conditions of (s1) and (s2) in ordering semantics.

> (s1) $\mathbf{O}(f)$ is true iff among the accessible possible worlds there are some f-worlds that are closer to the deontic ideal than any $\neg f$-worlds.
>
> (s2) $\mathbf{O}(r)$ is true iff among the accessible possible worlds there are some r-worlds that are closer to the deontic ideal than any $\neg r$-worlds.

Since there are no $r \wedge f$-possible worlds accessible from i, (s1) and (s2) cannot be true at the same time, let alone $\mathbf{O}(f \wedge r)$. So although the predicament of Sartre's pupil seems to be possible, the set of three sentences that represent his predicament is inconsistent when rendered in ordering semantics.

As we have seen in Kratzer's solution, given $f(i) = \{r \rightarrow \neg f, f \rightarrow \neg r\} = \{\neg(f \wedge r)\}$ and $g(i) = \{r, f\}$, (s1) and (s2) do not hold at all. According to this solution, instead of (s1) and (s2), the following negations hold:

(s4) It is not the case that he *ought* to join the Free French forces. $\neg\mathbf{O}(f)$

(s5) It is not the case that he *ought* to remain with his mother. $\neg\mathbf{O}(r)$

(s4) and (s5) together say that Jean is not expected to do either. This result simply does not get the nature of moral dilemmas right. If Jean has neither obligation, then there is no dilemma; he has nothing to worry about!

2.4.2 The Conflict Account

Now I will turn to the *conflict account*, the first of two straightforward ways to remedy the problems just raised for the simple account, still working within the framework of ordering semantics. The basic idea of the conflict account is this:

when there are two conflicting moral principles that ask a moral agent S to do ϕ and ψ respectively, but the agent cannot realize both ϕ and ψ, the following two *ought*-statements hold: "S ought to do ϕ" and "S ought to do ψ." For these two *ought*-statements to hold at the same time, we can evaluate each *ought*-statement against a *sub*-ordering source: one of maximally consistent *subsets* of a contextually provided ordering source. This idea is formally captured in ordering semantics as follows:

> $[\![\mathbf{O}(\phi)]\!]^{g(i),f(i)} = 1$ at i iff $\exists g_n(i) : \{w : w \in \cap f(i) \wedge \neg\exists v \in \cap f(i) : v \prec_{g(i)} w\} \subseteq [\![\phi]\!]$, provided that $g_n(i)$ is one of the maximally consistent subsets of $g(i)$.
>
> $\mathbf{O}(\phi)$ holds against $\langle g(i), f(i)\rangle$ if and only if there is at least one maximally consistent subset of $g(i)$, say $g_n(i)$, given which some ϕ-worlds are deontically better than any $\neg\phi$-worlds.

Recall the modal base and ordering source of the unfortunate predicament of Sartre's pupil again: $f(i) = \{\neg(r \wedge f)\}$ and $g(i) = \{r, f\}$. Given the modal base $f(i) = \{\neg(r \wedge f)\}$, there are two maximally consistent subsets of the ordering $g(i)$: $g_1(i) = \{r\}$ and $g_2(i) = \{f\}$. $g_1(i)$ is a maximally consistent set of propositions and ranks f-worlds higher than $\neg f$-worlds, and $g_2(i)$ is a maximally consistent set of propositions and ranks r-world higher than $\neg r$-worlds. Then, (s1) is true given $g_1(i)$ and $f(i)$; (s1) is true give $g_2(i)$ and $f(i)$; (s3) is true for either $g_1(i)$ or $g_2(i)$ and $f(i)$. Both *ought*-statements, (s1) and (s2), hold in the pupil's dilemma situation:

$$[\![\mathbf{O}(f)]\!]^{f,g} = [\![\mathbf{O}(r)]\!]^{f,g} = [\![(r \rightarrow \neg f) \wedge (f \rightarrow \neg r)]\!]^{f,g} = 1 \ (g = g_1 \cup g_2).$$

According to the conflict account, there could be multiple maximally consistent sub-ordering sources, and seemingly conflicting *ought*-statements can hold tout court in the same context by virtue of there being different maximally consistent sub-ordering sources for that context. This approach can be motivated and justified by the following consideration as well. We take multiple roles in society, and some of our duties and obligations coming from our different roles can conflict. Sartre's pupil *ought* to join the forces given his patriotic duty, while he *ought* to remain with his mother given his filial duty. (s1) $\mathbf{O}(f)$ holds in light of his patriotic duty, while (s2) $\mathbf{O}(r)$ holds in light of his filial duty. For each of Jean's different roles, a different ordering source is in play for each *ought*-statement. This approach allows us to describe a dilemma situation using two conflicting *ought*-statements:

(s1) Jean *ought* to join the Free French forces. $\mathbf{O}(f)$

(s2) Jean *ought* to remain with his mother. $\mathbf{O}(r)$

Intuitively, Jean's situation is accurately described by the conjunction of (s1) and (s2). This is exactly what the simple account above fails to capture.

Despite these advantages of the conflict account, however, it has some drawbacks. First of all, the analysis of *ought* in the conflict account is a big departure from a standard understanding of necessity modals in ordering semantics. Ordering semantics provides a structurally isomorphic analysis for modals of different flavors. Is there any reason to give a deontic necessity modal like *ought* a different analysis? Traditionally, in ordering semantics the truth-condition of necessity modal claims of different flavors (alethic, epistemic, deontic etc,) are defined in terms of comparative betterness of a set of possible worlds that satisfy most propositions in the relevant ordering source $g(i)$. For example, "$\Box(\phi)$," holds when ϕ-worlds are ranked higher (or comparatively better) than $\neg p$-worlds in that p-worlds meet

more propositions in $g(i)$ than any $\neg p$-worlds do. According to the conflict account, however, $\mathbf{O}(\phi)$ holds for any ϕ in $g(i)$. Although this analysis allows us to talk about each normative principle accepted in a deontic ordering source and to express them with a corresponding *ought*-statement, it seems that this analysis loses its major feature as a necessity modal.

Second, it seems that the conflict account for *ought* fails to capture some central aspects of the meaning of *ought*-statements. In the conflict account, *ought*-statements describe certain relevant moral facts accurately even in a dilemma situation; e.g., they describe obligations of Jean's that emerge from his roles as a son and a citizen. However, in a sense this means that these *ought*-statements merely restate the moral principles that are already registered in an ordering source. And some might reasonably have a different view about how to use *ought*s. Thomson ([92]) expresses her concern about the attempt to allow conflicts in all-things-considered *ought*s based on conceptual considerations of the distinction between commitments and *ought*s.

> It should be stressed that what is an odd idea is that "I ought to give C a banana" and "I ought to give D a banana" are both true [even though I have only one banana to give, and cannot give it to both C and D]. There is no oddity in the idea that "I am committed to C to giving C a banana" and "I am committed to D to giving D a banana" are both true. Similarly for the ordinary English expressions "obligation" and "duty"([92]); 83).

In this passage, Thomson gets at the idea that what really can be in conflict are commitments, not *ought*-statements.

> Some philosophers have canvassed the idea in recent years that it can be the

> case that I ought to do alpha and ought to do beta despite the fact that I cannot do both alpha and beta. Should we agree with them? It is an odd idea. I will certainly feel you have been unhelpful if when I tell you about my predicament, and ask what I ought to do, you tell me? Well, as a matter of fact, you ought to give C a banana and you ought to give D a banana." I just told you I have only one banana ([92]; 83).

Thomson puts forward this idea rather impressionistically, but she has a point. The conflict account undermines the purpose of using and reasoning with normative language as an action-guiding enterprise. In other words, in moral reasoning it is standard to use *ought*-statements to ask the question, "What ought I to do?" and to answer this question. The meaning of *ought* given in the conflict account seems to miss this kind of meaning. In the conflict account, *ought*-statements accurately describe one's normative predicament based on one's duties and obligations, what is required for one to do, and one's commitments. They are moral principles. However, ordinarily it seems that the *ought*-statements that we derive from the moral principles are supposed to be the answers to the deliberative question, "what ought I to do?"; by contrast, in the conflict account, *ought*-statements merely reiterate relevant principles that are already registered in an ordering semantics.

Third, the conflict account faces some traditional problems that moral dilemmas pose to SDL. According to the conflict account, a pair of directly conflicting obligations holds whenever there are directly conflicting propositions in an ordering source. Witness $g(i) = \{p, \neg p\}$. Therefore, a widely accepted principle like (AND) must be rejected or restricted somehow.

The Aggregation Principle (AND)

$\vdash \mathbf{O}(\phi) \wedge \mathbf{O}(\psi)) \supset \mathbf{O}(\phi \wedge \psi)$.

Since (AND) is indispensable for certain aspects of our moral reasoning, and it permits us to infer what to do from different obligations, a reasonable response would be to restrict the principle.[11]

The demands of the conflict account do not stop here. In particular, the truth conditions of *ought*-statements under the conflict account invite complications with the notion of weak permission. Weak permission is often defined as the absence of a prohibition. Here are two widely accepted principles regarding the notion of weak permission and the connection between permission and *ought*-statements. The dual principle tells us that a proposition ϕ is permissible iff it is not the case that ϕ ought not to be the case.

Deontic Dual

$\mathbf{P}(\phi) \leftrightarrow \neg\mathbf{O}(\neg\phi)$

The dual principle is more plausible once we assume that an *ought*-statement expresses an obligation: doing ϕ is permitted iff there is no obligation to do $\neg\phi$. The entailment principle tells us that anything that ought to be the case must also be permissible.

Entailment of Permission

$\mathbf{O}(\phi) \supset \mathbf{P}(\phi)$

However, when we use the conflict account to try to understand permission based on these seemingly innocuous and intuitive principles, we get some problem-

[11] This point is strongly made by van Fraassen ([95]), Horty ([48]) and Goble ([32]). In particular, Goble ([32]) explores several possible ways of rejecting or revising the Aggregation Principle to accommodate moral dilemmas.

atic results. Based on this idea, we can give the truth condition for permission, "$\mathbf{P}(\phi)$", as follows:

$[\![\mathbf{P}(\phi)]\!]^{g(i),f(i)} = 1$ at i iff $[\![\mathbf{O}(\neg\phi)]\!]^{g(i),f(i)} = 0$ at i.

Given the conflict account, there are two possible ways of implementing the dual: either to give $\mathbf{P}(\phi)$ a parallel truth condition to $\mathbf{O}(\phi)$ as in the conflict account or to take the truth condition of $\mathbf{P}(\phi)$ as the negation of the truth condition of $\mathbf{O}(\neg\phi)$.

[Permission 1]

$[\![\mathbf{P}(\phi)]\!]^{g(i),f(i)} = 1$ at i iff for some $g_n(i)$, $[\![\mathbf{O}(\neg\phi)]\!]^{g_n(i),f(i)} = 0$, provided that $g_n(i)$ is a maximally consistent subset of $g(i)$.

[Permission 2]

$[\![\mathbf{P}(\phi)]\!]^{g(i),f(i)} = 1$ at i iff there is no $g_n(i)$, for which $[\![\mathbf{O}(\neg\phi)]\!]^{g_n(i),f(i)} = 1$, provided that $g_n(i)$ is a maximally consistent subset of $g(i)$.

Let's assume that we accept [Permission 1] and the conflict account for Jean's dilemma. Given two maximally consistent ordering sources, $g_1(i) = \{r\}$ and $g_2(i) = \{f\}$, in his situation. For $g_1(i)$, $\mathbf{O}(f)$ does not hold; for $g_2(i)$, $\mathbf{O}(r)$ does not hold.

$[\![\mathbf{O}(f)]\!]^{g_1(i),f(i)} = [\![\mathbf{O}(r)]\!]^{g_2(i),f(i)} = 0$

Given [Permission 1] , $\mathbf{P}(\neg f)$ and $\mathbf{P}(\neg r)$ holds in Jean's situation because their truth conditions are met for g(i), f(i) as follows:

$[\![\mathbf{P}(\neg f)]\!]^{g(i),f(i)} = 1$ at i iff for some $g_n(i)$, $[\![\mathbf{O}(f)]\!]^{g_n(i),f(i)} = 0$, provided that $g_n(i)$ is a maximally consistent subset of $g(i)$.

$[\![\mathbf{P}(\neg r)]\!]^{g(i),f(i)} = 1$ at i iff for some $g_n(i)$, $[\![\mathbf{O}(r)]\!]^{g_n(i),f(i)} = 0$, provided that $g_n(i)$ is a maximally consistent subset of $g(i)$.

As a result, Jean is permitted not to join the force and Jean is permitted not to stay with his ailing mother: $\mathbf{P}(\neg f)$ and $\mathbf{P}(\neg r)$. But it is absurd to infer from one's dilemma that failing to meet either obligation is permissible. What about [Permission 2]? It says, $\mathbf{P}(\phi)$ is permissible iff there is no maximally consistent ordering source for which $\mathbf{O}(\neg\phi)$ holds. [Permission 2] does not have the counterintuitive result about Jean's dilemma that [Permission 1] does. It has its own problem: for Jean, neither to join the force nor to stay with his mom is permitted. Jean's dilemma situation is characterized by the pair: $f(i) = \{\neg(r \wedge f\}\}$ and $g(i) = \{r, f\}$. Given the modal base $f(i) = \{\neg(r \wedge f)\}$ and the two maximally consistent subsets of the ordering source, $g_1(i) = \{r\}$ and $g_2(i) = \{f\}$, for $g_1(i)$, $\mathbf{O}(\neg f)$ holds and for $g_2(i)$, $\mathbf{O}(\neg r)$ holds. Therefore, neither $\mathbf{P}(r)$ nor $\mathbf{P}(f)$ holds in Jean's dilemma situation. This result is even more problematic in that it conflicts with the entailment principle, the standard assumption that obligation implies permission.[12] If Jean is not permitted to stay with his mother, then in what sense does he have the obligation to stay with his mother? If Jean is not permitted to go join the force, then in what sense does he have the obligation to go join the force? The original problem of the simple account is: neither $\mathbf{O}(f)$ nor $\mathbf{O}(r)$ is true in the dilemma situation. It seemed that this problem was addressed by using multiple ordering sources, since it allows us to say: "$\mathbf{O}(f)$ and $\mathbf{O}(r)$."

[12]Brink calls this "The Weak impermissibility principle": $\mathbf{O}(\phi)$ implies $\mathbf{P}(\phi)$. About this principle he says, "The new deontic principle is the weak impermissibility principle. But surely that must be true. If it's impermissible for me to torture my neighbor, then surely it's not the case that I'm not obliged to torture him." ([9], 236).

However, this approach was vulnerable to some technical complications once we introduced the notion of permission into the picture. Of course, the deontic dual principle is a more controversial principle than the dual between necessity and possibility. And, one might take this complication itself to constitute a good reason to reject both principles relating permission and obligation (or *ought*-statements). However, it seems that we certainly have the concept of weak permission defined in terms of the notion of obligation or *ought*-statements. So these complications with the notion of permission will remain to be explained in the conflict account.

In sum, the conflict account allows us to describe one's moral predicament in the most straightforward way, but it has a cheapening result: in the conflict account, *ought*-statements merely reiterate normative principles that are already accepted in an ordering source. And, they are not adequate for expressing moral judgements as answers to the question of what one ought to do. So they are not very helpful in our normative reasoning to figure out what to do. Aside from these conceptual considerations, the conflict account also has technical complications with widely accepted principles regarding conjunction and permission. For example, on the conflict account, the aggregation principle (AND) is not valid for *ought*-statements. It is a bit surprising that this most basic inference rule should be rejected or revised. Moreover, the new understanding of *ought*-statements in the conflict account requires some substantial modifications and reconsideration of some commonly accepted principles concerning the relation between obligations and permissions. It seems that the conflict account gives us something we want regarding moral dilemmas, but requires too much besides: the price is too steep.

2.4.3 The Disjunctive Account

Turn now to the *Disjunctive Account*, the second straightforward way that ordering semantics might attempt to improve on the simple account. The basic idea of the disjunctive account is this: when there are two conflicting moral principles that ask a moral agent S to do A and B respectively, but the agent cannot realize both A and B, the following one *ought*-statement holds: "S ought to do A or B." This idea is formally captured in ordering semantics as follows:

> $[\![\mathbf{O}(\phi)]\!]^{g(i),f(i)} = 1$ at i iff $\forall g_{n(i)} : \{w : w \in \cap f(i) \wedge \neg\exists v \in \cap f(i) : v \prec_{g(i)} w\} \subseteq [\![\phi]\!]$, provided that $g_{n(i)}$ is one of the maximally consistent subsets of $g(i)$.

> $\mathbf{O}(\phi)$ holds against $\langle g(i), f(i)\rangle$ if and only if for all maximally consistent subset of $g(i)$, say $g_{n(i)}$, given which some ϕ- worlds are deontically better than any $\neg\phi$-worlds.

Again, for the modal base and ordering source of the predicament of Sartre's pupil, $f(i) = \{\neg(r \wedge f)\}$ and $g(i) = \{r, f\}$, there are two maximally consistent subsets of the ordering $g(i) : g_1(i) = \{r\}$ and $g_2(i) = \{f\}$. For each sub-ordering source, neither $\mathbf{O}(r)$ nor $\mathbf{O}(f)$ holds, but $\mathbf{O}(r \vee f)$ holds against each sub-ordering source. Therefore, the following disjunctive *ought*-statement holds for Sartre's pupil in this dilemma situation.

> (s6) Jean ought to join the Free French forces or remain with his mother. $\mathbf{O}(f \vee r)$

The disjunctive account is free of many of the unpalatable results that the

conflict account has. In the disjunctive account, it is easier to explain the action-guiding role of *ought*-statements. When there are two conflicting duties and obligations, even though we do not fulfill both, at least we should fulfil one or the other of them.

Nevertheless, this approach is not completely free of problems. One question is: even if there is no logical tension or inconsistency, can a dilemma situation really be simply defined as a circumstance in which we ought to make either one of the available choices? If it is a disjunctive *ought*-statement that tells us what to do in a moral dilemma situation, then our choice in the dilemma situation should not be so hard. For Jean, what he ought to do is either to join the forces or stay with his mother. In that case, Jean's choice should not be hard; he could just choose either one of them. He can fulfill the disjunctive *ought*-statement, (s6), by joining the force, and he can also fulfill (s6) by staying with his mom. By doing either one of them, he can easily fulfill his duty; there is no violation. This means that in the disjunctive account, we cannot express the violation of an obligation via a conjunction of $\mathbf{O}(\phi)$ and $\neg\phi$. For although Jean has an obligation to remain with his mother, if he refrains from doing so in order to liberate his country then no obligations were violated.

This result is highly counterintuitive from the perspective of moral psychology. The disjunctive interpretation treats the two courses of actions, joining the forces and staying with his mother, as equally good options with respect to all of the relevant moral considerations, the "overabundance of choice" situation. There is nothing wrong with choosing one option over the other. This solution avoids contradictions, but it seriously cheapens our anguish as moral agents in the face of genuine moral dilemmas. There is no reason for Jean to be so conflicted on

the day of his decision. If this disjunctive approach is right, then Jean's inevitable guilt after his choice is ungrounded—even irrational—because he has done nothing wrong.

Marcus ([67]) puts forward her argument for the existence of genuine moral dilemmas based on the existence of this moral residue and inevitable remorse and sense of guilt. In a dilemma situation, either choice of a moral agent will be followed by remorse or guilt. These feelings are appropriate only if the agent believes that she had failed to fulfill her duties or obligations. However, this very intuitive explanation does not work if we accept the disjunctive account. If, after deciding to leave his ailing mother to join the forces, Jean tries to explain why he feels bad by saying, "I *ought* to remain with my ailing mother, but I'm not going to," it is simply false, since neither (s1) nor (s2) holds; his sense of guilt is unexplained because there is no violation.[13]

Here I do not argue that a semantic account should explain this phenomenon of moral psychology. It seems that moral residue and regret are explained by the existence of violation or by describing the situation as a violation situation, and the violation of an obligation is stated most simply as the conjunction of an *ought*-statement and the negation of is prejacent: $\mathbf{O}(\phi) \wedge \neg\phi$. A semantic theory cannot *explain* a phenomenon in moral psychology, but it should be able to adequately represent what we think the true description of the phenomenon is; otherwise it can be disputed on the same philosophical grounds that motivate that description. If the disjunctive account is right, then Jean's predicament is no more serious than that of Buridan's ass, which cannot decide between identical hay stacks and thus starves to death. Absent substantive philosophical defense, this is a problem

[13]Some might wondering if this is an independent problem of the disjunctive account, something in the vicinity of free choice permission. $\mathbf{O}(\phi \vee \psi)$ sounds like a free choice permission in the natural language reading of a disjunctive "ought," unlike "$\mathbf{O}(\phi)$ or $\mathbf{O}(\psi)$."

for the disjunctive account. Jean's hesitation and difficulty in making a choice in his predicament should not be degraded into a meaningless agony in front of equivalent options. The most serious problem of this approach is that it cheapens the difficulty of the decision in a dilemma situation.

2.4.4 The Dilemma for Ordering Semantics

Given the consideration of moral dilemmas, which account of *ought*-statements should we take? Which account better captures the meaning of *ought*-statement in our normative discourse? The conflict account allows us to most adequately describe a dilemma situation as one of two true, conflicting *ought*-statements. But this account is technically demanding and requires some substantive restrictions on commonly accepted principles: namely, Aggregation and the principles regarding permission and obligations. In general, this truth condition of *ought*-statement departs from ordering semantics' standard account of necessity modals. And at a conceptual level, the meaning of *ought*-statements as represented in the conflict account differs in important ways from the meaning of *ought*-statements that are used for moral judgments.

On the other hand, the disjunctive account gives a better analysis for *ought*-statements as moral judgments. This account allows us to advise a person in a dilemma situation that she ought to do either ϕ or ψ, even though neither choice is desirable and there is no plausible way of breaking this tie. However, this account results in essentially equating dilemma situations with "overabundance of choice" situations. Given the characterization of a moral dilemma in the disjunctive account, Sophie's choice is no more serious than that of Buridan's ass. Moreover, this framework has no resources to explain the moral residue that dilemmas leave

on us regardless of which choices we make. The disjunctive account is a more natural account of the meaning of deontic ought statements as deontic necessity modal claims. But it does not leave room for adequately presenting a conflicting pair of *ought*-statements.

Now the proponents of ordering semantics face their own dilemma between two ways of characterizing moral dilemmas in ordering semantics.

2.5 The Third Way: "Why Not Both?"

In the previous section, we have evaluated two possible ordering semantic accounts of *ought*-statements, in search of an adequate representation of our reasoning with conflicting *ought*-statements in dilemma situations. Each account comes with some costs, but it seems that they complement each other. Sometimes a nave question opens us to a new possibility: "Can't we have both?" In this section I am going to sketch what it is like to take this simple third option and what it takes to have both kinds of *ought*s in the ordering semantic framework.

The conflict account of *ought*-statements enables us to use *ought*-statements to talk about relevant normative principles that are accepted in an ordering source. Given the conflict account, for any proposition in an ordering source, the corresponding *ought*-statement holds. For example, given $g(i) = \{r, f\}$, both $\mathbf{O}(r)$ and $\mathbf{O}(f)$ hold. Even for $g(i) = \{b, \neg b\}$, both $\mathbf{O}(b)$ and $\mathbf{O}(\neg b)$ hold. Therefore, *ought*-statements as defined in the conflict account directly expresses accepted normative principles in an ordering source.

On the other hand, the truth-condition of *ought*-statements in the disjunctive

account capture evaluative judgments about what is good all-things-considered in a given circumstance. For example, given $g = \{a, b, c\}$ and $f = \{\neg a\}$, $\mathbf{O}(b \wedge c)$ holds. This kind of *ought*-statement tells us what is the most desirable thing to do all-things-considered in a given situation. Therefore, even in a dilemma situation, this type of *ought*-statement guides one's action since a disjunctive *ought*-statement still holds in a dilemma situation.

It is not surprising that there are different types of *ought*-statements, relating to different normative considerations that operate in our normative discourse and reasoning. A true deontic *ought*-statement can relate to different types of normative considerations. Consider an example of a simple *ought*-statement, "Sue ought to keep her promise". This *ought*-statement holds in two different situations: when it is a good thing for Sue to keep her promise in a given circumstance, and when there are certain duties or obligations that require Sue to keep her promise regardless of her current circumstances and even regardless of whether it would be good, all things considered, for her to keep the promise. When used in the former situation, this *ought*-statement expresses the speaker's evaluative judgment about what is the most desirable thing for Sue to do (all-things-considered) in her current circumstance. When used in the latter situation, the same *ought*-statement expresses the speaker's deontic judgment of what is required for Sue to do, as dictated by her relevant duties or obligations. Let's call the former kind of *ought*-statements axiological *ought*-statements, and call the latter kind deontological *ought*-statements. The axiological and deontological distinction relates to two different types of normative reasons: what is good all-things-considered in a given circumstance and what is required by a set of relevant duties and principles, respectively. This distinction in normative reasons is reflected in two distinct patterns of reasoning that we can observe in our reasoning and discourse that involve

ought-statements.

Of course, this distinction between two kinds of *ought*-statements will not come as news to moral philosophers. What I have been emphasizing is none other than the familiar, well-understood distinction between *pro tanto* obligations and all-things-considered obligations. My point is that even this familiar, basic distinction has so far eluded semanticists who are trying to represent moral reasoning in a formal way. Clearly, in ordering semantics, the meaning of deontological *ought*-statements is captured in the conflict account, whereas the meaning of axiological *ought*-statements is captured in the disjunctive account. But there is a need for a single account of normative language to capture both kinds of statements.

Thankfully, an at least rudimentary solution to this problem is easy to find. Based on the truth-conditions given in the conflict and disjunctive accounts, we can define both deontological and axiological *ought*-statements in the ordering semantic framework by introducing two operators. Let's use "$ought_D$" ($\mathbf{O}_D$) for deontological *ought*-statements defined by the conflict account, and "$ought_A$" ($\mathbf{O}_A$) for axiological *ought*-statements defined by the disjunctive account. Here are the formal definitions again.

> $[\![\mathbf{O}_D(\phi)]\!]^{g(i),f(i)} = 1$ at i iff $\exists g_n(i) : \{w : w \in \cap f(i) \wedge \neg\exists v \in \cap f(i) : v \prec_{g(i)} w\} \subseteq [\![\phi]\!]$, provided that $g_{n(i)}$ is one of the maximally consistent subsets of $g(i)$.
>
> $[\![\mathbf{O}_A(\phi)]\!]^{g(i),f(i)} = 1$ at i iff $\forall g_{n(i)} : \{w : w \in \cap f(i) \wedge \neg\exists v \in \cap f(i) : v \prec_{g(i)} w\} \subseteq [\![\phi]\!]$, provided that $g_{n(i)}$ is one of the maximally consistent subsets of $g(i)$.

The advantage of having two separate operators in ordering semantics is obvious

in the case of moral dilemmas. Let's go back to the story of Sartre's pupil, Jean. Jean's patriotic duty to join the forces and fight for his country and his duty as a son to take care of his mother are what are generally required of him, and thus it is only natural to understand them in terms of deontological *ought*-statements. Therefore, the most adequate description of his moral predicament is the conjunction of the two competing obligations binding him, and his practical ability:

(s7)	Jean $ought_D$ to join the Free French forces.	$\mathbf{O}_D(f)$
(s8)	Jean $ought_D$ to remain with his mother.	$\mathbf{O}_D(r)$
(s3)	If he joins the forces, he does not remain with his mother, and if he remains with his mother, he does not join the forces.	$(r \rightarrow \neg f) \wedge (f \rightarrow \neg r)$

On the other hand, an adequate piece of advice for him about what to do in the situation would be the following disjunctive axiological *ought*-statement.

(s6)	Jean $ought_A$ to join the Free French forces or remain with his mother.	$\mathbf{O}_A(f \vee r)$

Note that (s6-8) are compatible even when (s3) hold; in other words, they all hold for the same contextual parameters: $f(i) = \{\neg(r \wedge f)\}$ and $g(i) = \{r, f\}$.

The superiority of this third account compared to the conflict account and the disjunctive account is clear and obvious. First, Jean's unfortunate predicament is adequately described with the conjunction of deontological *ought*-statements : $\mathbf{O}_D(f) \wedge \mathbf{O}_D(r)$. He would be advised and his action should be guided by the disjunctive axiological *ought*-statement: $\mathbf{O}_A(f \vee r)$.

Second, this two-*ought* strategy distinguishes moral dilemma situations from "overabundance of choices" situations. All things considered, Jean $ought_A$ to join

the forces or stay with his mom: $\mathbf{O}_A(f \vee r)$. Likewise, we can make a similar disjunctive claim about Buridan's ass: all-things-considered, Buridan's ass *ought*$_A$ to eat the haystack or to drink the water in the bucket: $\mathbf{O}_A(h \vee w)$. However, the existence of corresponding normative principles distinguishes these two choice situations in a meaningful way. Jean is forced to choose between fighting for his country and taking care of his ailing mother because of the two conflicting duties he has; thus, two corresponding deontological *ought*-statements hold in this dilemma situation: $\mathbf{O}_D(f) \wedge \mathbf{O}_D(r)$. On the other hand, for Buridan's ass, there are no true corresponding deontological *ought*-statements for its two possible choices.[14]

Third, this two-*ought* account has an easy and straightforward formula for violation: "$\mathbf{O}_D(\phi) \wedge \neg\phi$." This formal representation of the violation of a duty is possible because we can still refer back to *ought*-statements that have contributed to the ordering source. For most formal systems for deontic modality do not provide adequate formal representations of unfulfilled or violated *ought*s and unfeasible *ought*s (in a relevant sense). Fourth, we can explain why either choice Jean makes would lead to his deep compunction and regret for violating one of the duties: we explain this by pointing out the relevant deontological *ought*-statement—e.g., Jean ought to stay with his mother and he didn't. Moreover, we can avoid the cheapening result that we found in the disjunctive account. When Jean makes a choice, his remorse and guilt should be explained in terms of the violation of one of his deontological *ought*-statement. And, this moral, emotional residue is another key feature of a dilemma that the "overabundance of choices" situation does not

[14]Some might have a following question about the definition of moral dilemmas: if a moral dilemma is defined simply in terms of a pair of conflicting deontological *ought*-statements, wouldn't that mean that a moral dilemma is epidemic? I do not think this will be the result of the definition, because there is a disjunctive condition. There is a moral dilemma only if there is a pair of conflicting deontological *ought*-statements and the corresponding disjunctive axiological *ought*-statement holds because they are completely on par. There could be many pairs of conflicting deontological *ought*-statements. However, it must be not that common that two deontological *ought*-statements are completely on par or incommensurable.

exhibit. Horty ([48]) makes a point resonating with these last two advantages. He points out that unfulfilled and violated duties and obligations have significant explanatory roles in our moral reasoning and practice.

> . . . a defeated prima facie *ought* may nevertheless justify feelings of compunction or regret, or even generate certain reparational *ought*s, such as the need to explain or apologize: in the present example, after missing a lunch date with my friend in order to rescue the child, I would certainly feel the need at least to explain the matter to my friend, and perhaps to apologize as well, in a way that I would not if we had never even had a lunch date. The problem of characterizing the moral force carried by defeated prima facie *ought*s is, however, a complicated matter that I will not try to address here ([48]; fn.5)

The proposed two-*ought* approach can provide an adequate formal presentation of violated or defeated duties and obligations. This presentation respects the data from our moral psychology and actual moral practice: in particular, feelings such as regret, guilt and moral compunction that we have even when we know that we could not help but make the choice that we made.

In the previous section, we discussed two ways of semantically analyzing *ought*-statements in ordering semantics, and found that both accounts have their own plausibility and weaknesses, but they can be complementary to each other. I propose to have both types of *ought*-statements in a single semantic account. Some might wonder why there have to be two *ought* operators in our semantic account, when there is only one orthographic type, "ought." This is a fair concern. But note that different types of *ought*s and their different functions in normative reasoning are well recognized in moral theories. For example, Ross ([82])'s distinction

between all-things-considered *ought* and prima facie *ought* is widely accepted. A prima facie *ought*-statement expresses a broad reasoning for actions and motivations, while all-things-considered *ought*s reflects the result of integrating and balancing various reasons.[15] Different kinds of normative reasons and the different roles that they have in our normative reasoning and judgments provide us a good reason to have more than one *ought* operator in a semantic account. In this paper, I have shown one way of having two types of *ought*s in the ordering semantic framework and illustrated how useful they are in adequately identifying and representing moral dilemmas.

2.6 Conclusion

In this paper, I discussed the challenges that formal accounts of normative language face regarding the possibility of moral dilemmas— in particular, ordering semantics for *ought*-statements. I discussed two possible ordering semantic accounts of *ought*-statements in order to accommodate the possibility of moral dilemmas without inconsistency: the conflict account and the disjunctive account. Each account gets some aspects of our moral reasoning right, but each has its own serious drawbacks. As I explore the pros and cons of the two accounts, the advantages of combining these two views become obvious. The importance of having various *ought*s that reflect different kinds of normative reasons and the different roles they have in our normative reasoning and judgments are obvious. To provide an adequate formal account for our normative reasoning in various situations, including

[15]It is now common to replace Ross's notion of prima facie duties with pro tanto reasons. Pro tanto reasons count in favor of or against some act as far as they go, but capable of being defeated by opposing reasons. For more discussion of various kinds of reasons and oughts, see Greenspan ([33]).

dilemma situations, a semantic account should be able to incorporate different notions of *ought*s as well. Finally, I have shown how the unified account can be executed in the ordering semantic framework and that there are many potential advantages of having two *ought*s in a formal account.

CHAPTER 3

MORAL TWIN EARTH AND GENUINE DISAGREEMENT

R.M. Hare ([38])'s Cannibals argument and Horgan and Timmons (H&T, hereafter)' Moral Twin Earth argument are two famous arguments in metaethics that are designed to disprove the descriptivist account of the meaning of moral terms.[1] According to this view, moral language works exactly like descriptive language in that moral terms refer, and their reference is determined in the way that the reference of descriptive terms is determined. Both arguments deploy the following simple observation about genuine disagreements: to have a genuine disagreement between two parties, they have to mean the same thing by the words used in their dispute; otherwise they are just talking past each other. So let's call this type of argument the Disagreement-Based Argument (DBA). Despite its far-reaching influence and significance in metaethics, what this argument really assumes and how it works have not been carefully examined. In this paper, I analyze the structure of the DBA argument and identify its most problematic premise, and discuss how the proponents of descriptivism can and should respond to the DBA argument, in particular regarding the elicited intuitions from DBA thought experiments.

An instance of DBA invites us to imagine an alternative speech community whose terms such as 'good,' 'right' and 'wrong' play the same role in reasoning and practice as our terms 'good,' 'right' and 'wrong' do: 'right' is a term of praise, whereas 'wrong' is a term of criticism. But they apply those terms to different types of actions, persons, and practices from those to which we apply our terms. So for

[1]Hare (1952, 1989); Horgan and Timmons ([42], [43], [44], and [45]); and Gibbard ([28]) also raise a similar consideration against cognitivism. In particular, Horgan and Timmons present several versions of the Moral Twin Earth argument as a general recipe for arguments against the descriptivist account of moral terms. This account goes along with a form of moral realism: namely, the ethical naturalism endorsed by Boyd ([7]) and Brink ([9]), and the moral functionalism of Jackson ([49]).

some types of actions, persons, and practices, members of the alternative community would say, "it is (morally) right", while we would say, "it is not (morally) right." If the meanings of moral terms are their descriptive contents, then it follows that the members of this alternative community ascribe different properties with their words 'good,' 'right' and 'wrong' from what we ascribe with our words 'good,' 'right' and 'wrong.' Therefore, DBA concludes, descriptivism fails to explain these apparent genuine disagreements in morality between the alternative community and us.

First, I am going to introduce Hare's Cannibals argument and H&T's Moral Twin Earth argument, and their target view, descriptivism, and extract the common argumentative structure of DBA. I am going to identify four major premises of DBA and show how they support the argument's conclusion. In particular, I am going to focus on the link between there being a genuine disagreement between two disputing parties and these parties' agreeing in meaning, and identify the assumption that DBA makes about this link, which I call [Platitude]. I will provide some linguistic evidence to debunk [Platitude] and conclude that first-order semantic incompatibility is not a necessary condition for there to be genuine disagreement between two prima facie disputing parties. Without [Platitude], the argumentative force of DBA is seriously undermined. Even after rejecting [Platitude], our intuition that there are genuine moral disagreements between two separate communities in a DBA scenario will remain and will need to be explained. For it is not immediately clear how the proponents of descriptivism, in particular natural moral realists, would deal with the intuition that there is something substantial at issue between the two speech communities in a DBA scenario in the realist picture. There are two possible rejoinders that the proponents of descriptivism can make: either to accommodate this intuition at face value in one's theory by

providing an adequate notion of a genuine disagreement, or to explain it away by accounting for why we have such an intuition, without accepting what it says at face value. In what follows, I will critically examine Plunkett and Sundell's ([72]; P&S hereafter) recent strategy to accommodate the intuition by understanding the disagreement in a DBA thought experiment as what they call a metalinguistic negotiation. But it will turn out that P&S' account is seriously limited, as it cannot make sense of genuine moral disagreement between two causally isolated speech communities, as in H&T's Moral Twin Earth scenario. Finally, I put forward a rather unembellished story that the proponents of descriptivism can adopt to explain the disagreement intuition away—Moral Twin Earthlings (MTEs) and Earthlings (Es) disagree about what to do.[2] I claim that there is nothing to be ashamed of if they take this "explaining away" rejoinder. Consequently, there is greater wiggle room for descriptivists than it appears at first. DBAs raise some interesting challenges to the proponents of descriptivism, but do not necessarily disprove descriptivism: Rather, the crucial step of DBA, [Platitude], is not a stable principle, and the proponents of descriptivism can explain why we tend to believe that there are genuine moral disagreements between two parties in DBA scenarios in terms of disagreement about what to do and how to live. Therefore, it is hasty to conclude that descriptivism is false.

[2]The seed of this approach is found in Copp ([19]) and Merli ([68]). My approach is different from theirs in that I present this account to debunk the disagreement intuition after defusing the original challenge against descriptivism. Also, I will provide a further development of this idea that avoids some immediate problems to it. One of those problems will be about whether cognitivists can take this strategy without giving up on their descriptivist account of moral language.

3.1 Disagreement-Based Arguments and Descriptivism

According to traditional metaethical descriptivism, moral terms refer, and their reference is determined in the way that the reference of certain non-normative terms (e.g., natural kind terms) is determined. The general spirit of this semantic account is naturally endorsed by natural moral realists. According to them, moral thought and discourse purport to describe some truth about moral properties which supervenes on or is identical with some facts about metaphysically unproblematic properties. Different versions of realism provide different stories about normative ontology and their versions of descriptivist semantics. Against this normative semantics, Hare ([38]) presents a powerful argument as follows.

> Let us suppose that a missionary, armed with a grammar book, lands on a cannibal island. The vocabulary of his grammar book gives him the equivalent, in the cannibals' language, of the English word 'good'. Let us suppose that, by a strange coincidence, the word is 'good'. And, let us suppose, also, that it really is equivalent— that it is, as the Oxford English Dictionary puts it, 'the most general adjective of commendation'... If the missionary has mastered his vocabulary, he can, so long as he uses the word evaluatively and not descriptively, communicate with them quite happily. They know that when he uses the word he is commending the person or object that he applies it to. The only thing they find odd is that he applies it to such unexpected people, people who are meek and gentle and do not collect large quantities of scalps; whereas they themselves are accustomed to commend people who are bold and burly and collect more scalps than the average.
>
> We thus have a situation that would appear paradoxical to someone who thought 'good' ...was a quality word like 'red'. Even if the qualities in people

> that the missionary commended had nothing in common with the qualities that the cannibals commended, yet they would both mean what the word 'good' meant. If good were like 'red', this would be impossible; for the cannibals' word and the English word would not be synonymous. It is because in its primary evaluative meaning 'good' means neither of these things, but is in both languages the most general adjective of commendation, that the missionary can use it to teach the cannibals Christian morals. ([38]; 146-9)

Michael Smith underlines the argumentative force of Hare's Cannibals argument as follows:

> ...if the cannibals use their words 'good' and 'right' to refer to the causes of their uses of the word 'good' and 'right,' and the missionaries use their words 'good' and right' to refer to the causes of their uses of the words 'good' and 'right', and if no more can be said about the content of their respective judgments, then radical relativism is on the horizon. For we seem to have good reason to suppose that the causes of the cannibals' and the missionaries' uses of the words 'good' and 'right' are very different from each other. And in that case we cannot suppose that the cannibals and the missionaries disagree with each other about what is really good and right. ([85]; 34)

Hare and Smith conclude that the apparent disagreement between the cannibals and the missionary does not make sense if descriptivism is true; therefore, descriptivism is false. Horgan and Timmons revive this argument with their Moral Twin Earth scenario and deploy various versions of it to argue against versions of descriptivist semantics endorsed by realists: Causal Regulation Semantics developed by Cornell Realists Boyd and Brink, and Jackson's Neo-Descriptivism for

his moral functionalism.[3] In fact, the general argumentative strategy employed by Hare and H&T applies to any normative semantics that provides a reference fixing mechanism for moral terms that has nothing to do with their normative roles. According to any such normative semantics, two moral terms simply have different meanings if they are applied to different kinds of actions and practices, even if they play the same kind of &role in discourse and practice (e.g., they both play roles of commendation).[4]

Now turn to H&T Moral Twin Earth scenario and their version of DBA, which targets Boyd's Causal Regulation Semantics. Here is their original setup of the Moral Twin Earth scenario.

> Now consider Moral Twin Earth [...] is just about like good old Earth: same geography and natural surroundings [...] people who live in the twin United States by and large speak Twin English. [...] Of particular importance here is the fact Moral Twin Earthlings have a vocabulary that works much like human moral vocabulary; they use the terms "good" and "bad," "right" and "wrong" to evaluate actions, persons, institutions and so forth [...] Let us suppose that investigation into Twin English moral discourse and associated practice reveals that their uses of twin moral terms are causally regulated by certain natural properties distinct from those that regulate English moral discourse. The properties tracked by twin English moral terms are also func-

[3]H&T directly raise a challenge of explaining our intuitions about the MTE scenario in H&T ([42]), and they put pressure on descriptivism by reviving G.E. Moore's Open Question Argument ([69]) in H&T ([43]) and J.L. Mackie's Argument from Queerness ([66]) in H&T ([44]) with the MTE thought experiment. Here I am focusing on the version they present against descriptivism which the original target of Hare ([38])'s DBA argument.

[4]For the causal regulation semantics for naturalistic moral realism, see Boyd ([7]), Brink ([9]) and Railton ([78], [79], [80]); for neo-descriptivism see Jackson ([50]); Copp ([19])'s realist-expressivism is another metanormative account that does not rule out the Moral Twin Earth possibility.

> tional properties, whose essence is functionally characterizable by means of a normative moral theory. But these are non-consequentialist moral properties, whose functional essence is captured by some specific deontological theory; call this theory T_d. [...] In addition, suppose that if Twin Earthlings were to employ in a proper and thorough manner the same reliable method of moral inquiry which (as we are already supposing) would lead Earthlings to discover that Earthling uses of moral terms are causally regulated by functional properties whose essence is captured by the consequentialist normative theory T_c, then this method would lead the Twin Earthlings to discover that their own uses of moral terms are causally regulated by functional properties whose essence is captured by the deontological theory T_d. ([43]; 164-5)

Orthographically identical twin moral terms and moral terms play the same roles in moral and twin moral discourses and practices, but have different extensions on each planet. For example, just as the orthographic type 'good' serves as a term of praise in English, so, too, does it serve in Twin English. Similarly, for English speakers, applying 'good' to something commits them to doing it or to preferring it over other acts or practices; so, too, for 'good' in Twin English. However, twin morality and morality could be diverse enough that Moral Twin Earthlings would sincerely and truly utter and believe, "lying in order to save a person's life is wrong," while Earthlings would sincerely and truly utter "No, lying in order to save a person's life is not wrong." Even if both parties speak truly, however, most people react to this example by thinking that MTEs and Es disagree about a normative issue. H&T deploy this intuition as critical evidence to argue against descriptivism— à la Putnam's Twin Earth argument for externalism (against content-internalism).[5]

[5]Putnam ([77], [76]). H&T's Moral Twin Earth scenario is parallel to Putnam's Twin Earth

> On the one hand, we could say that [...] the moral terms used by Earthlings rigidly designate the natural properties that causally regulate their use on Earth, whereas the twin moral terms used by Twin Earthlings rigidly designate the distinct natural properties that causally regulate their use on Twin Earth; hence, moral and twin moral terms differ in meaning, and are not intertranslatable. On the other hand, we could say instead that moral and twin moral terms do not differ in meaning or reference, and hence that any apparent moral disagreements that might arise between Earthlings and Twin Earthlings would be genuine disagreements, i.e. disagreements in moral belief and in normative moral theory, rather than disagreements in meaning. We submit that by far the more natural mode of description, when one considers the Moral Twin Earth scenario, is the second. [...] But if CSN [Boyd's Causal Semantic Naturalism)] were true, and the moral terms in question rigidly designated those natural properties that causally regulate their use, then reflection on this scenario ought to generate intuitions analogous to those generated in Putnam's original Twin Earth scenario. That is, it should seem intuitively natural to say that here we have a difference in meaning, and that Twin English "moral" terms are not translatable by English moral terms. Yet when it comes to characterizing the differences between Earthlings and twin Earthlings on this matter, the natural-seeming thing to say is that the differences involve belief and theory, not meaning. ([43]; 165-6)

H&T conclude that this descriptivist account of the MTE scenario does not square

scenario in that the only difference between two planets is the extension of the term(s) in question. However, note that the MTE scenario is generally held to elicit an opposite intuition from that about Putnam's Twin Earth scenario: the twin moral terms and moral terms have the same meaning.

with our intuition that there are genuine moral disagreements between Moral Twin Earthling and Earthlings, and thus, descriptivism is false.

2. The Structure of the Disagreement-Based Argument

Both Hare and H&T's arguments have the same argumentative structure, which consists of four building blocks:

1. **Alternative**: It is possible for there to be a hypothetical alternative speech community that has terms that work much like our moral terms regarding their roles in reasoning, discourse and practice, but have different extensions from our moral terms.[6]

2. **Descriptivist Result**: What Descriptivists Have To Say about [Alternative]
The alternative community's moral terms and our moral terms do not mean the same thing since they are not co-extensive.

3. **Data**: Competent English speakers' Intuitions about [Alternative]
For some term F in our moral vocabulary, when the people from the alternative speech community truly and sincerely claim "x is F" and we truly and sincerely claim "x is not F," they *genuinely disagree* in moral belief and normative moral theory, not in meaning.

[6]For names and natural kind terms, the intensional contents or concepts of them do not determine their extensions. Likewise, according to the Boyd-style Causal Regulation Semantics, the extension of a moral term is not determined by its intensional content or concept. So this account allows the possibility that the very same moral concepts have different extensions. Also, the crucial element in the MTE thought experiment is that the inhabitants of Moral Twin Earth and Earthlings are distinct linguistic communities in that the subsets of their vocabularies that include moral terms are causally regulated by different natural properties, despite the sameness of their roles in moral and twin moral discourses and practice. This assumes that the extension which causally regulates a term's use is not determined by the term's role in discourse and practice.

4. **Platitude**: The Link Between Genuine Disagreement and Agreement in Meaning

Two parties have a genuine moral disagreement with each other only if (i) they mean the same thing by the linguistic expressions used in the dispute, and (ii) their claims and beliefs at issue have semantically incompatible contents.

The argumentative force created by the MTE scenario is straightforwardly analyzed with these four elements in the form of *Modus Tollens* as follows:

[**Alternative**] Suppose that Moral Twin Earthlings' twin moral term 'right' has the same role in twin moral discourse, reasoning, and practice as Earthlings' moral term 'right' has, but the referent of the twin moral term 'right' on Moral Twin Earth is different from that of the moral term 'right' on Earth.

1. [**Descriptivist Result**] If descriptivism is true, then the twin moral term 'right' on Moral Twin Earth and the moral term 'right' on Earth have different meanings.

2. [**Platitude**] If the twin moral term 'right' on Moral Twin Earth and the moral term 'right' on Earth have different meanings, then the apparent moral disagreement between a Moral Twin Earthling's utterance "x is right" and an Earthling's utterance "x is not right" is not a genuine moral disagreement.

3. [**Data**] The disagreement between Twin Earthlings and Earthlings is a genuine moral disagreement.

4. Therefore, descriptivism is not true.

Both Hare's and H&T's arguments as instances of DBA show that the relevant

[Data] and [Descriptivist Result] about [Alternative] are in conflict given [Platitude] and conclude that descriptivism is false.

The structure of DBA is straightforward and each step seems to be plausible. However, the four elements of DBA are not completely free of concerns, and descriptivists may raise challenges to each. Given the descriptivist accounts of moral terms, [Alternative] is possible in principle. However, the possibility of [Alternative] can be challenged depending on which version of normative realism one combines with one's descriptivism. For example, some natural moral realists might insist that when the similarities between twin moral and moral reasoning and practices are sufficient for us to think that twin moral and moral terms have the same roles, twin moral terms refer to the same properties as moral terms do, or at least their extensions should not be so different that they have to be characterized by different moral theories.[7]

Even when [Alternative] is possible, [Data] can be challenged in various ways. H&T assume that our intuitions about the MTE scenario as competent English speakers are reliable data for semantic theorizing, like our intuitions about Twin Earth in Putnam's thought experiment. Dowell (2012) challenges this comparison, as well as the probative value of our intuitions about the MTE scenario. [Data] might turn out not to be the conclusive or reliable evidence for semantic theorizing. Nevertheless, we should be unsatisfied with the idea of dismissing this intuition as mere linguistic confusion. At least, we should explain why we have intuitions such

[7]This approach has its own plausibility, and natural moral realists like Boyd and Brink could claim that regardless of what moral theories prevail on each planet, the meaning of moral terms purely depends on what properties of actions, policies, characters, and other evaluable phenomena causally regulate the use of moral terms on each of the two planets. And, if the natural properties that regulate the use of Twin Moral Terms are different enough from those regulating the use of moral terms on Earth that they have to be characterized by different moral theories and twin moral and moral judgements diverge on important normative cases, it is not obvious whether we have [Data] on [Alternative]. For this type of response see Merli ([68]).

as [Data] about those scenarios.

[Platitude] is the most critical piece in DBA and is responsible for generating the direct tension between [Data] and [Descriptivist Result]. The notion of 'genuine disagreement' itself and the connection between genuine disagreement and agreement in meaning are not very clear. So [Platitude] itself deserves a careful examination more than any other elements of DBA. In the next section, I am going to carefully examine [Platitude] and challenge the plausibility of this core step in DBA.

3.2 Does Genuine Disagreement Require Agreement in Meaning?

What counts as a genuine or substantive disagreement and what accounts as a mere verbal dispute are substantial issues in philosophy that deserve serious discussion.[8] I am not going to delve into these questions directly, but I'm going to get clear on what kind of characterizations of genuine disagreement are required for DBA and whether the notion of genuine disagreement is stable enough to ground DBA. Both Hare's and H&T's arguments exploit a certain understanding of the nature of genuine disagreement to draw their intended conclusions. The basic idea of their understanding of genuine disagreement can be captured with this impressionistic slogan: "Genuine moral disagreement requires agreement in meaning." In the previous section, I have specified this idea in terms of two necessary conditions in [Platitude], which plays a key role in DBA by connecting genuine disagreements

[8]See Chalmers ([16]) and Jenkins ([51]) for the attempts to define what counts as a mere verbal dispute. See MacFarlane ([64], [65]) for an overview of various notions of disagreements.

and the ways of using words (or the meaning of the words used) in a dispute.

[Platitude] Two parties have a genuine moral disagreement with each other only if (i) they mean the same thing by the linguistic expressions used in the dispute, and (ii) their claims and beliefs at issue have semantically incompatible contents.

Let's call the first necessary condition the same meaning condition and the second the semantic compatibility condition. These two semantic conditions adequately rule out some paradigm examples of mere verbal disputes from counting as genuine disagreements. For example, suppose that Joey says, "Cindy is tall", and Beth says, "No, Cindy is not tall." The disagreement between Joey and Beth would evaporate as soon as Joey and Beth learned that they have used different standards for applying the word "tall" to Cindy. For example, Beth meant that Cindy is not tall for a basketball player, while Joey meant that Cindy is tall for a 30-year-old woman. Moreover, what they say by making those utterances does not conflict in content. They can reach the agreement that Cindy is tall for a 30 year old woman, but not tall for a basketball player, without changing their minds. However, at least sometimes, a genuine disagreement does seem to be due to more than just the incompatibility of semantic contents expressed by two parties.

First, satisfying the same meaning condition is not necessary for genuine disagreement. Note, for example, that the semantic incompatibility condition does not require the same meaning condition. Suppose that Andy says, "Joon is marsupial, and Ben says, "No, Joon is not mammal." They disagree, and their claims are semantically incompatible. But such a semantic incompatibility does not require the words 'mammal' and 'marsupial' to have the same meaning. Even if there is a disagreement as semantic incompatibility, it is not sufficient to conclude that the relevant key terms must have the same meaning. This incompatibility or conflict

in content does not necessarily require that all the words used in a dispute be used co-extensively.

Now the question is whether genuine disagreement between two parties always requires the semantic incompatibility of their claims. Sundell ([90]) and P&S ([72]) argue that it does not. Consider the following short exchange.

> Anne: There is one proton in the nucleus of a helium atom.
>
> Ruth: No, there are two protons in the nucleus of a helium atom.

The propositions that these two speakers express are logically consistent. Yet, P&S argue that in this dispute the speakers really do disagree with each other, both intuitively and by the light of the incompatibility of the pragmatic contents of what they communicate. The two speakers in this dispute genuinely disagree because one communicates her belief that there is exactly one proton in the nucleus of a helium atom, while the other communicate her belief that there are exactly two protons in the nucleus of a helium atom. What this example shows is that for there to be genuine disagreement, the incompatibility does not have to be in what is literally expressed. There could a genuine disagreement due to the incompatibility of what they communicate beyond literal meaning. Moreover, there is nothing non-genuine about the disagreement, and indeed it is a disagreement very much worth having. The discussion so far shows that the two semantic conditions, (i) the same meaning condition and (ii) the semantic incompatibility condition, are not necessary conditions for two parties to have a genuine disagreement. And, our intuitions about what counts as genuine, meaningful, or substantive disagreements do not solely rely on what is literally expressed. Here is another interesting type of linguistic phenomenon that proves that even when two parties neither use

their terms in the same way nor express incompatible contents, there could be a meaningful disagreement between two. Recently, the metalinguistic usage of linguistic expressions as a type of linguistic phenomenon has been a subject of lively discussion by linguists and philosophers of language. ([5]) A linguistic expression is used metalinguistically when it is used to communicate information about the appropriate usage of that very expression in the relevant context. With this idea of metalinguistic usage, P&S ([72]) claim that we can even dispute which concept should be associated with a term by engaging in metalinguistic uses of that very term. They call such a dispute a metalinguistic negotiation.

It seems that metalinguistic usage is indispensable in our communication. It seems that we often have to address normative questions about how to use our words and which concepts we should employ in a particular linguistic practice. These are questions in conceptual ethics (Burgess and Plunkett, [13], [14]; B&P hereafter). Here's Barker's famous example that nicely illustrates a metalinguistic use of 'tall.'

> Normally, [3] will be used in order to add to the common ground new information concerning Feynman's height:
>
> [3] Feynman is tall.
>
> But [3] has another mode of use. Imagine that we are at a party. Perhaps Feynman stands before us a short distance away, drinking punch and thinking about dancing; in any case, the exact degree to which Feynman is tall is common knowledge. You ask me what counts as tall in my country. "Well," I say, "around here, ..." and I continue by uttering [3]. This is not a descriptive use in the usual sense. I have not provided any new information about the world, or at least no new information about Feynman's height. In fact,

> assuming that tall means roughly 'having a maximal degree of height greater than a certain contextually supplied standard', I haven't even provided you with any new information about the truth conditions of the word tall. All I have done is given you guidance concerning what the prevailing relevant standard for tallness happens to be in our community; in particular, that standard must be no greater than Feynman's maximal degree of height ([5]; 1-2).

In this context, the participants of the conversation try to find contextually appropriate usage of 'tall'—in Barker's terminology, a sharpening of the term 'tall'. In this case, it is assumed that there is a unique most accurate way of using the term 'tall' in the context. Thus, it is possible that we can express disagreement about what is the most contextually salient threshold for 'tall' by engaging in metalinguistic use of the descriptive term 'tall.' For instance, in the above situation, another person from the speaker's country can challenge the speaker's remark by saying, "No, Feynman is not tall." Again, this person also uses 'tall' metalinguistically to make a claim about what is the most contextually salient threshold for 'tall.' So, the disagreement is a factual one about which of two or more competing characterizations of the shared conversational context is most accurate. This could be true even if the word 'tall' has different extensions in the mouths of the two different speakers. Disagreements about conceptual ethics arguably can be found with normative terms as well. Consider the following recurring example in Chalmers ([16]), Sundell ([91]) and P&S ([72]). Suppose that in the context of a political debate, two speakers disagree about whether waterboarding is torture because they have different standards of what counts as torture.[9]

[9]To fill out the details, suppose further that Minsoo follows the United Nations in defining torture as any act inflicting severe suffering, physical or mental, in order to obtain information or to punish (United Nations 1984, 85), while Hyerin follows former U.S. Justice Department

Minsoo : Waterboarding is torture.

Hyerin : No, waterboarding is not torture.

The disagreement between these two speakers is a genuine and substantial disagreement in the context of discussing and deciding what is permitted and what is not during interrogation. They use different standards of what counts as torture and apply the word 'torture' to different sets of action-types. So at least some philosophers would say that Minsoo and Hyerin do not agree in the meaning of the key term, 'torture', and express different properties by the predicate 'is torture'. However, even if that is the case, nonetheless the debate that Minsoo and Hyerin engage in is not a mere verbal dispute, because learning about how they use their words does not remove the disagreement. On this reading, the reason that they have this exchange is not because of mere linguistic confusion; rather, which concept to associate with the word 'torture' is a substantial issue worth discussing. Their own views about how to use the word 'torture' are exactly what s/he ultimately wants to get across and what s/he disagrees about with the other.

These examples and our discussion so far show that the first order semantic features like using co-extensive terms or expressing semantically incompatible contents are not necessary for two parties to be engaged in a genuine disagreement, or for observers to judge that there is a genuine disagreement—although meeting both the same meaning condition and the semantic incompatible condition would be a sufficient condition for having a genuine disagreement. So the kind of link between genuine disagreement and sameness of meaning for moral terms that DBA

practice in defining torture as any such act inflicting pain rising to the level of death, organ failure, or the permanent impairment of a significant body function (U.S. Department of Justice 2002, 340A).

assumes is either spurious or too weak to rely on in DBA's *Modus Tollens*.

Disproving [Platitude] significantly undermines the argumentative force of DBA. In particular, without [Platitude] DBA does not pose a direct threat to a descriptivist account of moral language.[10] However, [Data] remains and puts some independent pressure on the proponents of descriptivism who endorse metaethical realism.

3.3 What We Can Say About Moral Twin Earth

Without [Platitude], DBA loses its initial argumentative force. However, it seems that [Data] itself raises an interesting challenge to the proponents of descriptivism. [Data] says that there is a genuine disagreement between us and the community described in [Alternative]. Interestingly, it is not immediately clear what realist descriptivists would say about [Data]. There are two routes descriptivists can take here: either (i) accommodate [Data] by providing an adequate notion of a genuine disagreement that descriptivists can accept, or (ii) explain [Data] away by explaining why we have this intuition, even though the intuition is not correct.

P&S ([72]) is a good example of the former strategy. P&S argue that with the metalinguistic usage of a moral term, two parties disagree about how to use the term, and the views that the two parties express and accept in the metalinguistic negotiation are in conflict in a traditional sense: conflict in content. This is a very interesting approach. Unfortunately, however, it will turn out that P&S's account cannot deal with all instances of DBA, contrary to P&S's claims.

[10]Note that the challenge to descriptivism would not be that descriptivism entails that there is no genuine disagreement in the Cannibal and MTE cases, but rather that descriptivism has no good way to account for people's intuition that there is genuine disagreement.

I am going to present a simple unembellished account of why there are genuine disagreements between the cannibals and the missionary and between MTEs and Es in terms of disagreement *about what to do*. The proposed approach will provide a homogeneous account of what makes the disagreements in the DBA scenarios genuine and meaningful. However, this account does not guarantee that the disagreement about what to do in a DBA scenario is a moral disagreement. Nevertheless, we can explain why we tend to believe these disagreements are genuine moral disagreements in the DBA scenarios. Therefore, this approach will count as the "explaining away" approach, unlike P&S's attempt.

This debunking strategy might not be satisfactory for those who believe that [Data] is the evidence for the existence of a substantial *moral* issue between us and the alternative speech community described in [Alternative]. However, I do not see any strong reason for realists to accept the disagreement intuition about the MTE scenario in [Data]. Why should we believe that ordinary people's intuition about the MTE scenario faithfully describes the true structure of moral reality across different possible worlds? Moreover, it seems that there is more to lose than gain when accepting [Data].[11] Still with the notion of disagreement about what to do we can provide a fairly plausible explanation of why we have the intuition that there is some substantial disagreement between MTEs and Es. I only hope that this is good enough for many readers.

[11]And, it is unclear what realist picture should be like in order to embrace the possibility. Accepting [Data] at face value would also invite the ineffability problem into the realist territory, which is introduced and discussed in Eklund ([20], [21]).

3.3.1 Metalinguistic Negotiation and Genuine Disagreement

To explain a genuine disagreement as a metalinguistic negotiation, P&S maintain a rather traditional understanding of a genuine disagreement, and apply it to communicated contents that go beyond first-order semantic contents: that is, to what is (pragmatically or metalinguistically) communicated in a given context by utterances and what is accepted by participants in the conversation, including plans and attitudes.[12] Here's their characterization of genuine disagreement.

> **Disagreement Requires Conflict in Content (DRCC)**
>
> If two subjects A and B disagree with each other, then there are some objects p and q (propositions, plans, etc.) such that A accepts p and B accepts q, and p is such that the demands placed on a subject in virtue of accepting it are rationally incompatible with the demands placed on a subject in virtue of accepting q. (Perhaps, though not necessarily, in virtue of q entailing not-p.)

P&S claim that in the light of DRCC, the two speakers in the waterboarding example genuinely disagree with each other because their views about how to use the word 'torture' conflict in content.[13]

[12]Here P&S provide a liberal understanding of contents which can be explicitly or implicitly communicated by various types of linguistic (or non-linguistic) activities. Of course, I can intend to convey the value I accept by exclaiming "We All Are Immigrants!" and so can somebody someplace in the Midwest holding his cardboard picket saying, "BUILD THE WALL." But it seems that the different values I and such a person may or may not successfully communicate are not the meanings of linguistic expressions or utterances explained in semantics. I worry that P&S may relax the concept of contents too much in order to endorse the rather traditional understanding of disagreement in terms of semantic incompatibility.

[13]Recall the waterboarding case. In the context of discussing how to treat prisoners, the two speakers in the example communicate their views about the common issue by using the word 'torture' metalinguistically. In the exchange, by the utterance Minsoo communicates the view that the word 'torture' should be used in a way that includes waterboarding as its instance, while by the utterance Hyerin communicates her view that the word 'torture' should be used in a way that excludes waterboarding from the extension of 'torture' in the given context. P&S point out that with these metalinguistic uses of the term they ultimately communicate the two

P&S apply this strategy to the Moral Twin Earth Argument to make sense of the genuine disagreement between MTEs and Es ([72]; 19-22). Their explanation is two-fold: the disagreement between MTEs and Es concerns conceptual ethics *before they meet*, and concerns conceptual ethics expressed *via* metalinguistic negotiation *once they meet*. It is straightforward to apply their strategy to the latter case, a hypothetical MTE scenario where MTEs and Es meet. Suppose that Bob from Moral Twin Earth and Chris from Earth meet up and talk about whether it is (morally) right to lie in order to save a person's life. They engage in a metalinguistic negotiation over which concept to use for the term '(morally) right'. In this hypothetical situation, the Moral Twin Earthling Bob and Earthling Chris share a context and have the common issue of what concept to use for '(morally) right'— that is, which concept to use to figure out what to do, when lying is the only thing one can do to save one's friend's life. Bob claims that the term 'right' should be used in a way that includes lying to save a person's life, while Chris claims that the term 'right' should be used in a way that excludes such an action from the extension of 'right'. So, the contents that Bob and Chris communicate in this hypothetical situation are in conflict and meet the DRCC criteria. Therefore, it's a genuine disagreement by P&S' lights.

In the original MTE scenario, however, MTEs and Es never meet or talk, and they are even causally isolated from each other. P&S claim that even when there is no shared context between two parties, they can (genuinely) disagree about conceptual ethics.

> On our proposal, Bob and Chris's disagreement concerns which concept to express with the term 'morally right'. Our proposal is that they each advo-

semantically incompatible views that waterboarding is problematic, and that waterboarding is not problematic.

cate a view about which concept is best suited to play a certain functional role in thought and practice, a role that includes matters of how to treat others, what to hold each other responsible for doing, and how to live more generally. One reason that Bob and Chris want their preferred concept to be the one expressed by the term 'morally right' is because they each believe (correctly) that whatever is called 'morally right' will likely play this role, and that people who think that 'morally right' should mean something different have a different view about what concept should play this functional role. [...]

Given this sort of disagreement in conceptual ethics, it is entirely sensible to suppose that before they engage in any linguistic exchange — before they are even aware of one another's existence, much less one another's language — Bob and Chris have views about which concept should play this important functional role in thought and practice. They would have those views entirely independently of any awareness of the other person, or of the existence of another language. Bob's view is that the analytically utilitarian concept C1 should play this functional role in organizing our lives and Chris's view is that the analytically Kantian concept C2 should play that role. These views are incompatible, irrespective of whether Bob and Chris engage in a conversation where they express these views. Thus, by the lights of DRCC, they disagree, irrespective of whether they engage in conversation. In turn, if they do start talking to each other it is entirely sensible that they have a disagreement not just about which concept to deploy, but which concept to deploy using the expression 'morally right'. Thus, the normative issue becomes not only which concept to use, but also which concept will be paired with this particular word 'morally right'. That linguistic exchange, once it

happens, can proceed via metalinguistic negotiation. ([72]; 22)

When the two speech communities are causally isolated from each other, they do not really communicate in a single context. P&S claim that nevertheless the view that MTEs accept about conceptual ethics and the view that Es accept about conceptual ethics *conflict in content*.

However, this strategy does not work. Even if we accept P&S' liberal understanding of contents and what is communicated and expressed as including beliefs, plans and even views about conceptual ethics, P&S' analysis of the MTE scenario does not meet their DRCC criteria. Here's why.

According to P&S, Bob and Chris each advocate a view about which concept is best suited to play a certain functional role in thought and practice, a role that includes matters of how to treat others, what to hold each other responsible for doing, and how to live more generally. This convenient characterization of the important functional role leaves two possible ways of reading it: either it is the most general (thinnest) normative function of commending, praising, objecting or blaming in thought and practice, or it is a more specific kind of normative function that moral vocabulary has in our discourse and practice, such as *morally* commending, praising, objecting or blaming, or, on Moral Twin Earth, *twin-morally* commending, praising, objecting, or blaming. In P&S's argument, it is unclear whether they mean by this important function a general normative function or a moral function.

If P&S mean that Bob and Chris advocate different views about which concept should play the most general normative functional role in thought and practice, then it would turn out that their different views about conceptual ethics also cover the how to use "should." In short, they have different views about how to use

"should," so the occurrences of "should" in their views about conceptual ethics have different meanings.

> Bob: Utilitarian concept C1 *should* play an important functional role (the most general normative functional role) in thought and practice.
> Chris: Kantian concept C2 *should** play an important functional role (the most general normative functional role) in thought and practice.

Since they use the central word "should" differently, the contents of their views about conceptual ethics are not in conflict. Therefore, it does not meet DRCC. Now suppose that by the important functional role in thought and discourse P&S mean the function played by our moral concepts: *morally* commanding, praising, objecting and blaming—or, on Moral Twin Earth, *twin-morally* commanding, praising, objecting, etc. Note that there are different realms of normative considerations and practices, and that normativity comes not only from moral considerations but also from other evaluative perspectives like etiquette, economics, prudence and so on. It is certainly not that Bob from MTE and Chris from M have different views about which concept should play an important role in *economic* consideration and practice to figure out what to do and how to live as an *economic* agent. On Earth what we are considering is what concept we should use to *morally* praise, condemn, etc., whereas on Moral Twin Earth what we are considering is which concept we should use to *twin-morally* praise, condemn, etc. By hypothesis, morally praising/condemning is distinct from twin-morally praising/condemning. Then, Bob has a view about conceptual ethics in the domain of twin morality, while Chris has a view about conceptual ethics in the domain of morality.

> Bob: Utilitarian concept C1 should play an important functional role in *twin*

moral thought and practice to figure out what to do and how to live.

Chris: Kantian concept C2 should play an important functional role in *moral* thought and practice to figure out what to do and how to live.

Note that H&T's original scenario assumes that both MTEs are right about their twin moral judgments and Es are right about their own moral judgments since both twin moral terms and moral terms have *the right relationship* with relevant homeostatic clusters of twin natural and natural properties respectively—which causally regulate their uses of twin moral and moral terms. So the important functional role in MTEs' twin moral thought and practice is hooked up in the right way to a certain homeostatic cluster of twin natural properties, while the important functional role in Es' moral thought and practice is hooked up in the right way to a certain homeostatic cluster of twin natural properties, After elaborating what the important function is and the domains of thought and practice in Bob's and Chris's views about conceptual ethics, we can see that in either case they do not genuinely disagree with each other either because their thinnest normative concepts do not coincide with each other or because they talk about different domains of thought and practice. Bob's and Chris' views in conceptual ethics simply do not meet the DRCC criteria, and there is no conflict in content. For the situations where Bob and Chris meet or where the cannibals and the missionary meet, P&S's account deploying the notion of metalinguistic negotiation and DRCC would work. Not only the two parties in a meeting scenario engage in a meaningful dispute in order to coordinate their uses of the term at issue, but also the views they accept about how to use the common key term conflict in content in a traditional sense. However, this account cannot work for two speech communities causally isolated as in the MTE scenario, because they simply do not talk about the same thing.

3.3.2 Disagreement about What to Do and How To Live

Unfortunately, P&S' account of how genuine moral disagreement manifests in the MTE thought experiment is not satisfactory. However, accommodating [Data] is not the only option available for descriptivists, and might not be a smart thing to do. I don't see why descriptivists need to accept [Data] at face value and take up the burden of spelling out the mechanism of having genuine moral disagreements between causally isolated speech communities and even across possible worlds.[14] In this section, I am going to explore the second possible rejoinder to DBA, in particular the Moral Twin Earth Argument, that descriptivists actually can make. To explain [Data] away, what descriptivists should do is to explain why we have such intuitions about [Alternative]. The good news is that descriptivists can explain [Data] without altering their view about moral language. The explanation I propose is rather simple: Moral Twin Earthlings and Earthlings in the MTE scenario disagree about what to do in the situations where we naturally believe that we need to make moral judgements and actions. The intuition that H&T and Hare share is that, because descriptivists say that the participants in [Alternative] scenarios use their moral terms differently, the descriptivist account of moral terms cannot make sense of the existence of a real issue between them, or something that the disputants need to talk about and solve together. Once we conclude that Twin Earthlings use twin-moral terms while Earthlings use moral terms, the Twin Earthlings' twin moral judgment expressed by "A is (morally) right" and the Earthlings' moral judgment expressed by "A is not (morally) right" are not really in conflict: the semantic contents (that is, propositions) of what they believe and assert are compatible. However, the divergent uses of language inherent in twin

[14]Maybe some realists who endorse some version of objectivism or absolutism have a strong reason to accept [Data].

moral language and moral language lead to the divergence of twin-moral discourse and practices from moral discourse and practices. In other words, Moral Twin Earthlings and Earthlings could commit themselves to doing practically incompatible acts in some situation of the same type—in particular, when Moral Twin Earthlings truly and sincerely utter "A is (morally) right" and Earthlings' truly and sincerely utter "not-A is (morally) right" in the same situation. Therefore, the genuine conflict that we sense in Hare's Cannibals scenario and H&T's Moral Twin Earth scenario can be explained by pointing to the different actions they would take in a given situation: in other words, they disagree about what to do and how to live.[15]

Suppose that Oscar the Earthling and Toscar the Moral Twin Earthling read about the Trolley problem for the first time on each planet. Given H&T's MTE scenario, we would be able to see some disagreement very similar to the disagreement between utilitarians and deontologists. Oscar has a true belief that it is (morally) right to push the man on the bridge to save 17 people on Earth, while Toscar has a true belief that it is not (morally) right* to push the man on the bridge to save 17 people on Moral Twin Earth. And, each of them is willing to express their belief as follows:

Oscar : It is *right* to sacrifice one person's life to save 17 people.

Toscar : It is not *right* to sacrifice one person's life to save 17 people.

Despite appearances, their utterances are compatible, since 'right' in the Earthling's remark and 'right' in the Twin Earthling's utterance have different contents.

[15]Again, a disagreement based on the semantic incompatibility is not the only kind of genuine disagreement. The incompatibility of the semantic contents of what they believe and assert is not necessary for there to be genuine disagreement.

However, the genuine issue between them can be easily shown by asking them what they would do.

> Q: Suppose that you could save 17 people by pushing one person off the bridge in front of the trolley in the given situation. Would you push the person off the bridge?
>
> Oscar : Yes, I would sacrifice the man to save 17 people if I were in that situation.
>
> Toscar: No, I wouldn't sacrifice the man no matter what if I were in that situation.

Oscar and Toscar would commit themselves to doing, or praising, or commending different types of actions in the same hypothetical situation. These two types of actions cannot be taken by one single agent without changing her beliefs. What Oscar would do if he were in the Trolley situation is incompatible with what Toscar would do if he were in the Trolley situation, in that it is practically impossible for one to take Oscar's plan and Toscar's plan in the Trolley situation.[16]

Some might wonder if I am assuming too strong a connection between a moral belief and commitment or motivation that descriptivists as cognitivists cannot endorse. I do not think so. However, to avoid this worry we can appeal to a rationality or consistency in moral beliefs. As a rational (consistent) moral agent, Oscar is not supposed to *say* that he would not approve of sacrificing one man to save 17 people if he were in the situation after sincerely asserting, "it is *right* to sacrifice one person's life to save 17 people." What a rational moral agent believes and says are expected to be consistent at least. This consistency requirement

[16]Gibbard ([28], [29]) extensively discusses the logic of plans, including how disagreement in plan is possible. The notion of disagreement about what to do is indeed reminiscent of his disagreement in plan.

has nothing to do with whether they are motivated accordingly or have the right attitudes. Thus, this connection between what is said/believed and what would be said/believed by a rational moral agent is cognitivist-friendly.

Moreover, even stronger connection between what is said and believed by a moral agent and what that agent is committed to do in a hypothetical situation can be utilized by descriptivists here. It is simply because they are trying to explain why we tend to think that there is a substantial disagreement between MTEs and Es. As long as the connection between what Oscar says and believes and what he would be committed to do in a hypothetical situation is something ordinary English speakers routinely assume, descriptivists can utilize this assumed connection to explain their offhand intuition about the MTE scenario.

Some might argue that Oscar has a plan for what to do in the Trolley situation and Toscar has a plan for what to do if he were in the Trolley situation; hence, each of them has a different plan for themselves, and thus they are not in conflict. To understand this worry better, Suppose, for example, to get the best coffee in Ithaca Anne would go to Gimme Coffee, while Ruth would go to Ithaca Coffee. They can individually plan and get their favorite coffee from different places. Their plans are not in conflict. However, to the question of what would you do to get the best coffee in Ithaca, Anne and Ruth would come up with their different plans: the plan to go to Gimme Coffee and the plan to go to Ithaca Coffee. As an answer to the common question, their plans of what to do compete with each other. And, yes, whether or not two plans count as incompatible depends on what question the plans are an answer to. Likewise, as an answer to the Trolley situation question, their plans of what to do are competing answers, and they are in conflict in that one single agent cannot perform both plans in the situation at issue.

The discussion so far naturally suggests an alternative characterization of the necessary condition for a genuine disagreement.

> **[The Necessary Condition for Genuine Disagreement]**
> Two parties genuinely disagree with each other only if there is one single common issue to which they have competing responses.

We have the intuition that MTEs and Es genuinely disagree in morality because they would make different judgments and act differently if they were in the situations in which we tend to believe that moral judgements, decisions and acts are required. Pointing out the discrepancy in what they would do (or their plans) in some paradigm moral situations and the way they would live should be enough to explain why we have the intuition that there are genuine moral disagreements between them in the MTE scenario. That's it. This does not necessarily mean that there are inter-world moral issues or principles applying across possible worlds.[17]

Here I carefully limit the scope of the explanation. This "explaining away" strategy is not to explain what is exactly happening between MTEs and Es; rather, *to explain why we tend to take the immediate view* that MTEs and Es have a genuine disagreement in morality. I have explained this intuition with a more liberally understood notion of a genuine disagreement. Regardless of their awareness of each other's existence, we can indeed talk about the disagreement between them, when they have competing solutions, answers, responses, or beliefs about the common subject matter, whether or not it is factual or counterfactual. And, sometimes these competing answers are non-cotenable motivations, dispositions, plans and

[17] Also, we would simply not have a strong intuition that the disagreement between MTE and Es is a genuine moral disagreement if MTEs use 'ought' 'right' 'good' to praise and encourage some actions that we feel completely indifferent about—for example, MTEs claim, "harvesting green tomatoes is wrong," "stroking a kitten is (morally) right" and so on.

actions. And, this can explain the genuine disagreements in the DBA scenarios. And, we have the intuition that they are genuine disagreements in moral belief and theory because they look like our moral disagreements about what to do and how to live.

Some might find this "explaining away" approach is unsatisfactory, and think what [Data] says, "there is a genuine moral disagreement between MTEs and Es," has to be taken as face value. But I do not think that anyone has successfully explained [Data] at face value on behalf of descriptivists so far. Another reason I think the "explaining away" approach is the right one for descriptivists is this: they cannot say or admit that the disagreement between MTEs and Es is a genuine moral disagreement without specifying what they mean by "moral" there, given the fact that MTEs talk about their twin morality and Es talk about their morality. Other than saying that their judgments and practice fall into the category we tend to characterize as moral judgments and practice, what else can we say? In sum, once we reject [Platitude] and drop the belief that a genuine disagreement has to be characterized in terms of first-order semantic features, [Data] does not pose an unsolvable problem to descriptivism that is a semantic account for the meaning of moral terms. And, it is harmless even for realists who endorse descriptivism to point out non-cognitive states we usually associate with moral beliefs to explain why we naturally tend to believe that there is a genuine moral disagreement in a fanciful scenario like the MTE scenario.

3.4 Conclusion

So far we have discussed Hare's and H&T's versions of DBA that have been thought to pose a serious challenge to most versions of descriptivism that are accompanied by moral realism. I have identified the argumentative structure of these arguments and identified the key assumption of the argument: the link between a genuine disagreement and agreement in meaning ([Platitude]). This link has been widely accepted without much doubt. I have argued that the logical connection between a genuine disagreement and first-order semantic features is not strong enough to sustain DBA. Rejecting this link undermines the argumentative force of DBA significantly. However, [Data], our intuition that there is a genuine moral disagreement between two speech communities in [Alternative], would still remain to be explained. There are two types of rejoinders descriptivists can make regarding [Data]: either to accommodate [Data] by providing an adequate notion of disagreement, or to explain [Data] away. I have shown that one recent attempt to accommodate [Data] with the notion of metalinguistic negotiation is not successful. I have claimed that instead descriptivists can explain [Data] away by using an alternative concept of a genuine disagreement that goes beyond semantic incompatibility. This response does not make descriptivists and realists commit themselves to the claim that MTEs and Es really disagree about morality. I think that pure descriptivist realists cannot accept that there is a genuine moral disagreement between MTEs and Es—for such a disagreement would indeed have to be cashed out in terms of semantic incompatibility of what they literally express. However, at least they can mimic the expressivist narrative to explain the intuition away. For those who are skeptical about this strategy, I want to point out a well-known truism: our intuitions about fanciful thought experiments are not always reliable guidance of philosophical theorizing and not free of errors. Also, it is unclear what

it means to say that there is a genuine disagreement between MTEs and Es more than just its seeming existence. Without conclusive theoretical reason to do so, accepting [Data] at face value brings in unnecessary complications without many advantages.

BIBLIOGRAPHY

[1] Azizah Y Al-hibri. Deontic logic: A comprehensive appraisal and a new proposal. 1978.

[2] Lennart åQvist. Deontic logic. In *Handbook of Philosophical Logic*, pages 147–264. Springer, 2002.

[3] Ana Arregui. Chisholm's paradox in "should"-conditionals. In *Semantics and Linguistic Theory*, volume 18, pages 55–72, 2008.

[4] Ana Arregui. Detaching if-clauses from should. *Natural Language Semantics*, 18(3):241–293, 2010.

[5] Chris Barker. The dynamics of vagueness. *Linguistics and Philosophy*, 25(1):1–36, 2010.

[6] David I Beaver. *Presupposition and Assertion in Dynamic Semantics*, volume 16. Csli Publications Stanford, 2001.

[7] Richard N. Boyd. How to be a moral realist. In *Contemporary Materialism: A Reader*, page 307. Routledge, 1988.

[8] David O Brink. Moral conflict and its structure. *The Philosophical Review*, 103(2):215–247, 1994.

[9] David Owen Brink. *Moral Realism and the Foundation of Ethics*. Cambridge University Press, 1989.

[10] Aaron Bronfman and J. L. Dowell. the language of reasons and 'ought'. In Daniel Star, editor, *Oxford Handbook of Reasons*. 2018.

[11] Aaron Bronfman and Jl Dowell. Contextualism about deontic conditionals. *Deontic Modality*, page 117, 2016.

[12] John Broome. Reasons. In R. Jay Wallace, editor, *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*, pages 2004–28. Oxford University Press, 2004.

[13] Alexis Burgess and David Plunkett. Conceptual ethics i. *Philosophy Compass*, 8(12):1091–1101, 2013.

[14] David Plunkett Alexis Burgess. Conceptual ethics ii. *Philosophy Compass*, 8(12):1102–1110, 2013.

[15] Jennifer Carr. The if p, ought p problem. *Pacific Philosophical Quarterly*, 95(4):555–583, 2014.

[16] David Owen Chalmers. Verbal disputes. *Philosophical Review*, 120(4):515–566, 2011.

[17] Roderick M. Chisholm. Contrary-to-duty imperatives and deontic logic. *Analysis*, pages 33–36, 1963.

[18] Matthew Chrisman. On the meaning of 'ought'. *Oxford Studies in Metaethics*, 7:304–332, 2012.

[19] David Copp. Milk, honey, and the good life on moral twin earth. *Synthese*, 124(1-2):113–137, 2000.

[20] Matti Eklund. Alternative normative concepts. *Analytic Philosophy*, 53(2):139–157, 2012.

[21] Matti Eklund. *Choosing Normative Concepts*. Oxford University Press, 2017.

[22] Elizabeth G Epstein and Sarah Delgado. Understanding and addressing moral distress. *Online J Issues Nurs*, 15(3), 2010.

[23] Fred Feldman. *Doing the Best We Can: An Essay in Informal Deontic Logic*. D. Reidel Publishing Company, 1986.

[24] Stephen Finlay and Justin Snedegar. One ought too many. *Philosophy and Phenomenological Research*, 89(1):102–124, 2014.

[25] Dagfinn FøLlesdal and Risto Hilpinen, editors. *Deontic Logic: An Introduction*. Springer, 1971.

[26] Philippa Foot. Moral realism and moral dilemma. *The Journal of Philosophy*, 80(7):379–398, 1983.

[27] B. C. Van Fraassen. the logic of conditional obligation. *Journal of Philosophical Logic*, 1:417–438, 1972.

[28] Allan Gibbard. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Harvard University Press, 1992.

[29] Allan Gibbard. *Thinking How to Live*. Harvard University Press, 2003.

[30] Lou Goble. A logic of good, should, and would. *Journal of Philosophical Logic*, 19(2):169–199, 1990.

[31] Lou Goble. Murder most gentle: The paradox deepens. *Philosophical Studies*, 64(2):217–227, 1991.

[32] Lou Goble. A logic for deontic dilemmas. *Journal of Applied Logic*, 3(3-4):461–483, 2005.

[33] Patricia Greenspan. Practical reasons and moral 'ought'. *Oxford Studies in Metaethics*, 2:172–194, 2007.

[34] Jeroen Groenendijk and Martin Stokhof. Dynamic predicate logic. *Linguistics and Philosophy*, 14(1):39–100, 1991.

[35] Bengt Hansson. an analysis of some deontic logics. In FøLlesdal and Hilpinen [25], pages 121–147.

[36] Sven Ove Hansson. Deontic logic without misleading alethic analogies-part i. *Logique Et Analyse*, pages 337–353, 1988.

[37] Sven Ove Hansson. *The Structure of Values and Norms*. Cambridge University Press, 2001.

[38] Richard Mervyn Hare. *The Language of Morals*. Oxford University Press, 1952.

[39] Irene Heim. File change semantics and the familiarity theory of definiteness. 1983.

[40] Georg henrik Von wright. a new system of deontic logic. In *Danish Yearbook of Philosophy*, volume 1, pages 173–182. 1964.

[41] Georg henrik Von wright. a correction to a new system of deontic logic. In *Danish Yearbook of Philosophy*, volume 2, pages 103–107. 1965.

[42] Terence Horgan and Mark Timmons. New wave moral realism meets moral twin earth. *Journal of Philosophical Research*, 16:447–465, 1991.

[43] Terence Horgan and Mark Timmons. Troubles for new wave moral semantics: The "open question argument" revived. *Philosophical Papers*, 21:153–175, 1992.

[44] Terence Horgan and Mark Timmons. Troubles on moral twin earth: Moral queerness revived. *Synthese*, 92:221–260, 1992.

[45] Terence Horgan and Mark Timmons. Analytical moral functionalism meets moral twin earth. In Ian Ravenscroft, editor, *Minds, Ethics, and Conditionals: Themes from the Philosophy of Frank Jackson*, pages 221–236. 2009.

[46] John F Horty. Moral dilemmas and nonmonotonic logic. *Journal of Philosophical Logic*, 23(1):35–65, 1994.

[47] John F Horty. Nonmonotonic foundations for deontic logic. In *Defeasible Deontic Logic*, pages 17–44. Springer, 1997.

[48] John F Horty. Reasoning with moral conflicts. *NoûS*, 37(4):557–605, 2003.

[49] Frank Jackson. on the semantics and logic of obligation. *Mind*, 94(374):177–195, 1985.

[50] Frank Jackson. *from Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford University Press, 1998.

[51] Carrie Jenkins. Merely verbal disputes. *Erkenntnis*, 79:11–30, 2014.

[52] Hans Kamp. Discourse representation theory. *Natural Language at the Computer*, pages 84–111, 1988.

[53] Hans Kamp, Josef Van Genabith, and Uwe Reyle. Discourse representation theory. In *Handbook of Philosophical Logic*, pages 125–394. Springer, 2011.

[54] Niko Kolodny and John Macfarlane. Ifs and oughts. *The Journal of Philosophy*, 107(3):115–143, 2010.

[55] Angelika Kratzer. What 'must' and 'can' must and can mean. *Linguistics and Philosophy*, 1:337–355, 1977.

[56] Angelika Kratzer. the notional category of modality. In *Words, Worlds, and Contexts*, pages 38–74. 1981.

[57] Angelika Kratzer. Modality. In *Semantics: An International Handbook of Contemporary Research*, pages 639–650. 1991.

[58] Daniel Lassiter. Measurement and modality: The scalar basis of modal semantics. *New York: New York University Dissertation*, 2011.

[59] Edward John Lemmon. Moral dilemmas. *The Philosophical Review*, 71(2):139–158, 1962.

[60] David Lewis. *Counterfactuals*. Harvard U. Press, 1973.

[61] David Lewis. Semantic analyses for dyadic deontic logic. In *Logical Theory and Semantic Analysis*, pages 1–14. Springer, 1974.

[62] David Lewis. Ordering semantics and premise semantics for counterfactuals. *Journal of Philosophical Logic*, 10(2):217–234, 1981.

[63] Barry Loewer and Marvin Belzer. Dyadic deontic detachment. *Synthese*, 54(2):295–318, 1983.

[64] John Macfarlane. Relativism and disagreement. *Philosophical Studies*, 132:17–31, 2007.

[65] John Macfarlane. Varieties of disagreement. 2009.

[66] Jl Mackie. *Ethics: Inventing Right and Wrong*. New York: Penguin., 1977.

[67] Ruth Barcan Marcus. Moral dilemmas and consistency. *The Journal of Philosophy*, 77(3):121–136, 1980.

[68] David Merli. Return to moral twin earth. *Canadian Journal of Philosophy*, 32(2):207–240, 2002.

[69] George E Moore. *Principia Ethica:(1903)*, volume 960. Cambridge University Press, 1959.

[70] Peter L Mott. on chisholm's paradox. *Journal of Philosophical Logic*, 2(2):197–211, 1973.

[71] Paul M Pietroski. Prima facie obligations, ceteris paribus laws in moral theory. *Ethics*, 103(3):489–515, 1993.

[72] David Plunkett and Timothy Sundell. Disagreement and the semantics of normative and evaluative terms. *Philosophers' Imprint*, 13, 2013.

[73] Paul Portner. *Modality*, volume 1. Oxford University Press, 2009.

[74] Henry Prakken and Marek Sergot. Contrary-to-duty obligations. *Studia Logica*, 57(1):91–115, 1996.

[75] Henry Prakken and Marek Sergot. Dyadic deontic logic and contrary-to-duty obligations. In *Defeasible Deontic Logic*, pages 223–262. Springer, 1997.

[76] Hilary Putnam. *Reason, Truth, and History*. Cambridge University Press, 1981.

[77] Hilary Putnam. the meaning of meaning. In *Mind, Language, and Reality: Philosophical Papers Volume 2*, pages 215–271. Cambridge University Press, 2016.

[78] Peter Railton. Moral realism. *Philosophical Review*, 95(2):163–207, 1986.

[79] Peter Railton. Essentially general predicates. *Midwest Studies in Philosophy*, 18(1):166–176, 1993.

[80] Peter Railton. Noncognitivism about rationality: Benefits, costs, and an alternative. *Philosophical Issues*, 4:36–51, 1993.

[81] Joseph Raz. Reasons: Explanatory and normative. In *New Essays on the Explanation of Action*, pages 184–202. Springer, 2009.

[82] William D Ross. The right and the good. 1930.

[83] Alex Silk. Why 'ought' detaches: Or, why you ought to get with my friends (if you want to be my lover). *Philosophers' Imprint*, 14, 2014.

[84] Peter Singer. Famine, affluence, and morality. *Philosophy & Public Affairs*, pages 229–243, 1972.

[85] Michael R. Smith. *The Moral Problem*. Wiley-blackwell, 1994.

[86] Robert Stalnaker. Assertion. *Syntax and Semantics (New York Academic Press)*, 9:315–332, 1978.

[87] Robert C Stalnaker. A theory of conditionals. In *Ifs*, pages 41–55. Springer, 1968.

[88] William B Starr. Dynamic expressivism about deontic modality. *Deontic Modality*, pages 355–394, 2016.

[89] Christian StraßEr. A deontic logic framework allowing for factual detachment. In *Adaptive Logics for Defeasible Reasoning*, pages 297–333. Springer, 2014.

[90] Timothy Sundell. Disagreements about taste. *Philosophical Studies*, 155(2):267–288, 2011.

[91] Timothy Sundell. Disagreement, error, and an alternative to reference magnetism. *Australasian Journal of Philosophy*, 90:743–759, 2012.

[92] Judith Jarvis Thomson. *The Realm of Rights*. Harvard University Press, 1990.

[93] L Van Der Torre and Yao-hua Tan. An update semantics for deontic reasoning. *Norms, Logics and Information Systems*, pages 73–90, 1998.

[94] Jacobus Adrianus Van Eck. A system of temporally relative modal and deontic predicate logic and its philosophical applications. *Logique Et Analyse Louvain*, (99):249–290, 1982.

[95] Bas C Van Fraassen. Values and the heart's command. *The Journal of Philosophy*, pages 5–19, 1973.

[96] Frank Veltman. Defaults in update semantics. *Journal of Philosophical Logic*, 25:221–261, 1996.

[97] Kai Von Fintel and Anthony Gillies. An opinionated guide to epistemic modality. *Oxford Studies in Epistemology*, 2:32–62, 2007.

[98] Georg Henrik Von Wright. Deontic logic. *Mind*, 60(237):1–15, 1951.

[99] Malte Willer. Dynamics of epistemic modality. *Philosophical Review*, 122(1):45–92, 2013.

[100] Malte Willer. Dynamic thoughts on ifs and oughts. *Philosophers' Imprint*, 14:1–30, 2014.

[101] Bernard A.O. Williams and W.F. Atkinson. Symposium: Ethical consistency. *Proceedings of the Aristotelian Society*, pages 103–138, 1965.